

Fair Play Economics

Nobody Has to Lose.

A plain-language field guide to how money games are actually funded — and a set of worked designs for games, reward programs, and token systems in which no participant can end with less money than they started with

Douglas J. Brownfield

First edition — November 2026

Copyright

Fair Play Economics

Copyright © 2026 Douglas J. Brownfield

All rights reserved. No part of this book may be reproduced, stored, or transmitted in any form or by any means — electronic, mechanical, photocopying, recording, or otherwise — without the prior written permission of the author, except for brief quotations used in reviews.

This book was written with the assistance of AI tools, working from the author's own questions and editorial direction, and from a research record built for this book — the industry's own published datasets and funnels, firm disclosures, program documents, and the prior-art record assembled for each model.

The figures in this book are stated as of 2026, each with its basis printed where it appears. No audited industry-wide payout figure exists; where the record is open, the book leaves it open and says so.

Free editions of this book — EPUB, PDF, and read-online — are published by the author at Douglas J. Brownfield Books.

The four paintings in this book are by Douglas John Brownfield; they were made for FreeDough, the working arcade built on this book's models, and are reproduced here as plates. The twenty mechanics diagrams in Part III were drawn for this book from each model's own worked figures.

Contents

Introduction — The Only Question That Matters

Part I — How the Money Games Actually Work

Chapter 1 — The Funnel

Chapter 2 — Where the Payout Money Comes From

Chapter 3 — Breakage: The Quiet Harvest

Chapter 4 — Rake, Fees, and the Second Engine

Part II — Principles of No-Loss Design

Chapter 5 — Who Pays, and Why Paying Is Worth It

Chapter 6 — Transparency Is the Product

Chapter 7 — The Consideration Point

Chapter 8 — Name the Contribution

Chapter 9 — For the Builders, and the Players

Part III — The Models: Twenty Worked Designs

Chapter 10 — Funded by Marketing

Chapter 11 — Funded by Time

Chapter 12 — Funded by Attention and Data

Chapter 13 — Funded by People Who Want You Skilled

Chapter 14 — Funded by Each Other, Honestly

Chapter 15 — Funded by Savings and Surplus

Closing Case — The Saturday League

Conclusion — A Five-Minute Audit for Any Program

The Working Exhibits

Support

Back Matter

About the Art

About the Author

Also by Douglas J. Brownfield

Frontispiece



Plate — The Riverboat Herself

Art: Douglas John Brownfield

The four paintings in this book come from FreeDough, the working arcade built on the models in Part III, where they hang as the top shelf of its prize counter — the shelf reserved for the rewards a player earns only over a long season of play. They are reproduced here as plates because the tie between the book and the arcade is the point: this book sets out the designs, and

FreeDough is the same designs running live, where anyone can play them. The arcade is the book's working proof, and the book is the arcade's workings, shown.

Introduction — The Only Question That Matters

Every game that handles money answers one question, whether it prints the answer or not: *who pays?*

Ask it of a lottery and the answer is easy. The players pay. Their ticket money becomes the prize pool, the operator keeps a share, and the mathematics guarantees that the players, as a group, get back less than they put in. Nobody hides this. It is printed in the structure of the thing, and most players understand it perfectly well: they are buying a chance, and the chance is the product.

Ask it of a modern trading challenge — the kind sold online, where you pay a fee to prove you can trade, and in return you are promised a funded account and a split of the profits — and the answer takes some digging. The fee is a few hundred dollars. The funded account is often a simulated one: your trades are mirrored, not real, and no capital of the firm's is ever at risk on your skill. The payouts, when they come, are real money. So where does that money come from?

Mostly, from the next person's fee.

The numbers tell the story plainly. In the largest dataset available — more than 300,000 challenge accounts at ten firms, reported by FPFX, the industry's own technology provider — about 14 percent of accounts pass the evaluation and reach a funded stage. Of those, fewer than half ever receive a payout. Put together, roughly **7 percent of accounts ever reach a payout at all**. The average payout, when it comes, is around 4 percent of the account size the buyer was dreaming about.

Meanwhile the average buyer spends about $800, across about three attempts, chasing it. One caution belongs here, because this book's promise is that every number carries its basis: that 7 percent is a per-account figure from one vendor's sample. No audited industry-wide figure exists, and the firms that publish their own funnels report different results on different denominators — from about 2 percent per evaluation at one futures firm to about 17 percent per person at another. Chapter 1 sets those funnels side by side rather than averaging them into a fog. On every basis, though, the shape is the same: most of the people who pay in are never paid.

None of this requires a villain. It only requires arithmetic. If the prize pool is funded by the entrants, and the operator's costs and profit come out of the same pool, then the entrants as a group must lose. The only question the design leaves open is *which* entrants — and the honest name for that question is a lottery, whatever the marketing calls it.

This pattern — the participant quietly funding the prize — is not limited to trading challenges. It runs through a surprising share of the modern rewards economy. Loyalty points that expire unredeemed become income for the company that issued them; by the accounting literature's estimates, something like 15 to 40 percent of point value is never spent. Gift cards go partly unredeemed, year after year, at rates every retailer can predict. Tournament prizes are carved out of entry fees at a percentage that is rarely printed anywhere a player will see it. In each case the structure is the same: the people inside the system are the source of the money the system pays out, minus the house's share, and the details live in terms nobody reads.

This book is about the opposite design.

It turns out there are many ways to fund a game, a reward program, or a prize in which the participant is *never* the source of the prize. A merchant will pay for a verified customer walking through the door. A foundation will pay, year after year, out of the earnings on money it never spends — the Nobel Prizes have run that way for well over a century. An employer will pay for an employee's education because a trained workforce is worth more than the tuition. A city power utility will pay households to use less electricity at the expensive hour, because paying them is cheaper than firing up another power plant. A sponsor will post a prize for a prediction game simply because the attention is worth the price. In each of these systems, somebody pays — there is always somebody paying; money does not arrive from nowhere — but the payer is paying for something *they* want, and the participant's wallet is not the fuel.

That difference is the whole subject of this book. We will call a design **zero-loss** when no participant can end with less money than they started with. The term needs one immediate clarification, because zero-loss does not mean costless. People contribute real things to these systems: their attention, their time, their data, their skills, their behavior. A fair system does not pretend those contributions are nothing. It names them, openly, on the label — *your attention is the contribution here; here is what it earns; here is the share, printed in advance*. What a zero-loss system never does is reach into the participant's pocket for the prize money and call the reaching a game.

Part I of this book takes apart the systems most readers have already met — the challenge funnel, loyalty breakage, tournament rake — and shows, with their own published numbers, where the money actually flows. None of it is secret; it is merely scattered, and some of it is written in a dialect of disclosures designed to be skipped. Part II states the principles that separate a fair funding structure from a funnel: identify the payer and their real reason; print the rates; keep the prize pipeline under independent control; name the contribution. Part III is the constructive half: twenty worked designs for games and programs in which nobody has to lose. Each one names its payer, shows its arithmetic at a small, checkable scale, states what participants contribute, and points to the closest thing that already exists in the world — including, in several cases, the honest entry *none found*.

Two promises govern everything that follows. First, every claim that can be checked is sourced, and every number carries the basis it was measured on; where sources disagree, the disagreement is reported rather than smoothed over. Second, designs in this book are labeled for what they are: some are running in the world today at national scale, and some are original designs presented here for the first time. Those are different kinds of statements, and they are never blurred.

The reader this book is written for is two people at once: the builder deciding how to fund a program without feeding on its players, and the player deciding whether to enter a system whose funding they cannot yet see. Both need the same skill, and it is a small one. Before you join, before you build, before you believe — ask who pays, and keep asking until the answer is a name and a reason, not a fog.

If the answer is *you* — quietly, through fees, expirations, and fine print — you are not looking at a prize. You are looking at the pool.

Part I — How the Money Games Actually Work



Plate — The First Baker

Art: Douglas John Brownfield

Chapter 1 — The Funnel

The modern trading challenge is sold as a test of skill. Pay a fee, trade a simulated account against a set of rules, hit a profit target without breaking a loss limit, and the firm will "fund" you — give you a larger account and a split of the profits, commonly 80 percent to the trader. The industry grew from roughly ten firms five years ago to more than 400 selling challenges by 2025, by Bloomberg's count. To read the model, start with what the fee buys and what the funnel does to the people who pay it.

There are three billing shapes. The first is a one-time fee per attempt, standard at forex and CFD firms. FTMO, the market leader in Prague, lists its two-step challenge at €89 for a $10,000 account, €540 for $100,000, and €1,080 for $200,000; the targets are 10 percent in phase one and 5 percent in phase two, with a 5 percent maximum daily loss and a 10 percent maximum loss overall. Across 23 firms priced from their own live pages in September 2026, a roughly $100,000 program's headline price ran from $39.90 with a coupon to $969 at list — and the discounting is structural: list price is an anchor, and the coupon price is the price. The second shape is a monthly subscription until you pass, standard at futures firms. Topstep charges $49 a month for a $50,000 evaluation, then a one-time $149 activation fee for each funded account earned; the subscription converts attempts into tenure, and a slow learner is worth more than a fast one. The third shape is instant funding: pay a large premium — FTMO's instant $100,000 account costs €2,999 against €540 for the gated version — to skip the test entirely, under harsher trailing rules. The instant product is the cleanest proof available that the evaluation itself is a product, not a screening cost: the

same simulated account costs about five and a half times more without the gate.

The price points fill in the picture. FundedNext's Stellar two-step $100,000 program lists at $549.99, with a reset priced at $494.99 — about 90 percent of a fresh purchase, which makes the reset discount largely cosmetic at that tier. Subscription programs across the industry run $99 to $350 a month. MyForexFunds, historically, charged registration fees from $49 to $4,900 per account type, by the CFTC complaint's account of its records. And the discounting deserves its own sentence, because it teaches the reader how to hear every price in this industry: when one major futures firm runs coupons at 80 to 92 percent off as a standing practice, the list price is theater and the coupon price is the price. Figures in this chapter carry their basis as the book goes — audited filings where they exist (FTMO's, essentially alone), firm-reported disclosures where firms publish them, vendor-reported data where an industry supplier is the source, and alleged figures where a regulator's complaint is the source and the claim was never adjudicated. The labels are not decoration; in this industry, the basis is half the fact.

Now the funnel, and the denominators, which is where the industry's confusion lives. There is no audited industry-wide pass rate. The best public dataset is a vendor's: FPFX Technologies, a software provider to prop firms, reported on more than 300,000 accounts and roughly 100,000 traders at ten firms. In that dataset, about **14 percent** of accounts passed a challenge and received a funded account; about **45 percent** of funded traders ever reached a payout; together, about **7 percent of accounts ever reached a payout at all**. The average payout was about **4 percent of the nominal plan size** — roughly

$4,000 on a $100,000 plan. The average trader spent about $800 on challenge fees across about three challenges.

That figure needs its basis printed next to it, in this book, wherever it appears: in the largest published dataset — more than 300,000 accounts at ten firms, reported by FPFX Technologies, a software vendor to prop firms — about 7% of accounts ever reached a payout; no audited industry-wide figure exists, and firms that publish their own funnels report materially different results on different denominators, from about 2% per evaluation at Earn2Trade to about 17% per person at Topstep in 2025.

The firms that publish their own funnels report different things because they count different things, and the honest response is to print the funnels side by side rather than average them. Earn2Trade reported 8.89 percent of examinations passed in 2025, on a per-attempt basis. Topstep's 2025 disclosure reported 16.8 percent of evaluations completed — but 51.8 percent of *people* passed at least once, because buyers retry, and 33.3 percent of funded participants reached a payout, which works out to roughly 17 percent of people ever paid. Tradeify reported 17.2 percent of evaluation accounts passing — 40.3 percent of *people* passed at least once — and 28.5 percent of funded traders paid, roughly 11.5 percent of people. MyFundedFutures reported 19.97 percent of evaluations passed and 39.2 percent of funded traders paid — a second summary of the same disclosures says 28.56 percent, and both figures circulate; on its dated per-account basis the firm runs to roughly 12.4 percent of people ever paid, or about 5.8 percent per account. The rest of the published field shows the same spread: Take Profit Trader reported 36.22 percent of its shorter, easier tests passed;

HyroTrader 13.49 percent on its one-step and 8.70 percent on its two-step; FunderPro 7.35 percent of challenge buyers reaching funded; MyForexFunds' own February 2023 statistics claimed 19.05 percent passing phase one and 42.18 percent passing phase two. FTMO, FundedNext, and Apex currently publish no pass figures at all. CFD executives, speaking for the forex side of the industry, report roughly 1 to 2 percent of buyers ever paid. Two cautions complete the record rather than decorate it. No firm publishes a denominator-complete funnel — attempts to passes to funded to paid to live capital — except partially the futures firms above, so every "industry pass rate" in circulation is a sample with a viewpoint. And the largest firms' cumulative payout claims escalate by date in their own marketing — FTMO's public claims run from $450 million (2025) to "over $500 million" to "$650 million+" (2026), described by reviewers as an unaudited aggregate; no audited cumulative figure exists, and this book treats the three as dated snapshots, not a series. The pattern under the disagreement is consistent: per-attempt figures run low, per-person figures run higher because repeat buyers eventually pass at higher rates, and the funded stage is a second filter of similar strength — roughly 55 to 70 percent of funded traders never withdraw a dollar, losing the simulated account to drawdown and consistency rules first. One more figure completes the picture: at the futures firms that disclose it, between 0.71 and 4.2 percent of simulated-funded traders are ever moved to accounts trading real capital. "Live" capital is nearly fictional.

The rules deserve a plain reading, because they are the funnel's machinery. A daily loss limit, typically 5 percent, stops a bad day early. A maximum drawdown caps the total decline; in its

trailing form, used by most futures firms, the floor rises as the balance rises, so gains stop being safety and become a higher tightrope — a trader can be ahead overall and still be terminated by an ordinary pullback from a peak. Consistency rules, which cap the best day at 30 to 50 percent of total profit, prevent one fortunate day from carrying a target; one rulebook analysis estimates they can remove 33 to 50 percent of a single strong day's profit from payout eligibility. Minimum trading days force exposure across several sessions. An analysis of roughly 500,000 traders attributes about 70 percent of failures to hitting loss limits rather than missing profit targets. Whatever their fairness merits in any given case, each rule is, economically, a lever on the probability that a funded account ever becomes a payout — and payout gates continue after funding: waiting periods (FTMO's first reward is requestable on day 14), minimum profit thresholds, caps per payout cycle, and "winning days" requirements such as Topstep's five days of at least $150 profit. The complaint against MyForexFunds alleged the extreme version of the same machinery — drawdown findings used as pretexts and covert execution handicaps, delay and slippage profiles assigned to winning accounts — and the industry's internal correspondence, again as alleged, shows staff tracking "out-goings rising almost 100%" in volatile weeks and treating winning customers as a risk-management problem. Allegations are not findings, and Chapter 2 gives that case its full cautions; but the structural point survives them: a firm whose rules let it re-price a winner after the fact holds an option no printed rulebook discloses, and the option's value comes out of the same pool as the payouts.

Put the funnel together as unit economics, for 1,000 challengers on a blended $100,000-style plan. The assumptions are the FPFX figures above: average total spend of $800 per challenger across attempts, a 14 percent funded rate, 45 percent of funded traders ever paid, and average total payouts of about $5,600 per paid trader — an estimate built from the FPFX average of roughly $4,000 per payout and roughly 1.4 payout events before an account ends. Fee revenue is $800,000. Of the 1,000 challengers, 140 are funded and 63 are ever paid. Payouts cost $352,800 — 44.1 percent of revenue. Affiliate commissions, at a quarter of challengers arriving through affiliates paid 15 percent of fees, take about $30,000; payment processing at 3 percent takes $24,000. What remains before platform, support, and marketing costs is $393,200 — about 49 percent of everything the challengers paid in. The cross-checks land where they should: the 44 percent payout ratio matches the figure alleged for MyForexFunds in the CFTC's complaint, and sits just under the roughly 55 percent of revenue that FTMO's audited filings show going to payouts — FTMO pays a richer mix, including fee refunds to winners. The model is dominated by one product: the pass rate, times the paid rate, times the average payout. Every rule in the stack exists to keep that product small; every marketing claim exists to keep challengers from computing it.

Chapter 2 — Where the Payout Money Comes From

The finding that organizes this book is simple, and once seen it cannot be unseen: in the challenge model, the payouts are funded from the fee pool, and most of the "funded" trading never touches a market.

The funded account is simulated at essentially every major retail firm. FTMO states plainly in its product pages that its challenge, verification, and funded accounts are demo accounts with fictitious capital; the payouts — it calls them Rewards — are real money calculated on simulated profits. Topstep's evaluation and funded stages are likewise simulated. Even the firm's genuine market activity does not change the structure: FTMO says it executes real trades using the *data* of selected clients' simulated trades, which means the firm may copy-trade its best simulated traders with its own money while the trader is still paid on the simulated result. No capital of the firm's is at risk on any individual trader's skill. The money that pays winners has exactly one source of consequence: the fees of the challengers, current and past.

The extreme documented case is MyForexFunds. The CFTC's complaint, filed in August 2023, alleged that the firm collected at least $310 million in fees from more than 135,000 customers between November 2021 and August 2023, paid out about $137 million — roughly 44 percent of fees — and kept net income of about $172 million. The complaint alleged that the firm itself was the counterparty to substantially all trades, that fewer than 100 of its roughly 24,000 "live" customers ever had a single

trade routed to a real dealer, and that payments to traders came "from fees paid by other customers, in a manner similar to a Ponzi scheme."

Those figures are allegations, and their legal status must be stated as carefully as the figures themselves. The complaint's numbers were drawn from the firm's own records as the agency read them, and they were widely reported — but they were never adjudicated. On May 13, 2025, the district court dismissed the CFTC's case with prejudice, as a sanction, after a Special Master found the agency had acted willfully and in bad faith in its handling of evidence; the CFTC was ordered to pay $3,148,484.22 of the defendants' legal fees. The fraud claims were never proven and never rejected on the merits. Anyone citing "$310 million in fraud" as an established fact is overstating the record; anyone citing the dismissal as proof the model was clean is doing the same in reverse. What survives both cautions is the arithmetic, which does not depend on the case: payouts at about 44 percent of fees is the same ratio the worked example in Chapter 1 produces from the industry's own published funnel, and the same structure FTMO's audited filings show at about 55 percent.

FTMO is the respectable end of the same structure, and its filings are the more instructive document, because they show what the model does in a bad year. In 2024, on revenue of CZK 6.73 billion (about $323 million), FTMO earned a net profit of CZK 1.89 billion — a 28 percent margin. In 2025 revenue grew, to CZK 8.81 billion (about $423 million), but trader rewards came in about 82 percent above the reserve the firm had set aside: CZK 955 million paid against a CZK 523 million reserve, after a late-2025 rally in gold let crowds of traders pass and

profit by simply going long. Net profit fell 59 percent, to CZK 782 million — an 8.9 percent margin.

That is the hot cohort, and it is the model's structural weakness, not an accident of one year. The fee-pool model is stable only while winners are uncorrelated — one trader winning here, another there, each too small to matter. When a single market move makes the same trade right for thousands of challengers at once, pass rates and payout claims spike together, and the payout ratio can pass 100 percent of the fee stream. Run Chapter 1's 1,000 challengers with a trending market that doubles the pass rate to 28 percent, lifts the funded-to-paid rate to 55 percent, and raises the average payout to $7,000, and payouts become $1,078,000 — 135 percent of fee revenue, a loss of about $278,000 before operating costs. FTMO survived its version because it is large and cash-rich; it held CZK 2.52 billion in cash at the end of 2024. A thin firm in the same month simply does not pay, which is the industry's periodic collapse pattern: Finance Magnates estimated about 100 prop firms shut down in 2024 alone. The reserves, trailing drawdowns, consistency rules, and payout caps described in Chapter 1 are the four walls against that month. They are usually described to customers as trader discipline. In the firm's ledger, they are cost control, and both descriptions are true at once — which is precisely why a reader should insist on knowing the second one.

Around the fee pool sits a second engine of smaller revenues, and it matters because it shows how thoroughly the funnel monetizes every stage of a customer's passage through it. **Reset fees** are priced just below a fresh purchase — $59 to $169 at futures firms — to make continuing feel cheaper than restarting; one major futures firm abolished resets in favor of expiring

evaluations that must simply be rebought, the same mechanic with the discount removed, and funded-account resets can exceed evaluation prices (one firm's professional-tier reset is $649). **Activation fees** collect a second toll after success: Topstep's $149 per funded account, or roughly $59 to $69 per performance account elsewhere, multiplied across multi-account traders — a ten-account trader on promotional evaluations can pay about four times more in activation than in evaluations. **Add-ons** sell back the friction itself: split boosts from 80 toward 95 percent, drawdown boosts, payout-frequency upgrades. **Scaling plans** cost the firm little and retain the few who climb: FTMO adds 25 percent to the simulated balance per qualifying four-month block, raises the split to 90 percent, and caps accounts at $400,000 per strategy — the visible ladder keeps the whole cohort climbing toward a rung almost none will reach. **Platform and data fees** run quietly underneath (one futures firm charges $38 a month for Level 2 market data). **Affiliate commissions** — 8 to 20 percent of referred fees at FTMO — outsource customer acquisition to a review-and-discount-code ecosystem whose income depends on not emphasizing the payout rate. And above the individual firms sits a **white-label layer**: the infrastructure provider behind hundreds of branded storefronts keeps a spread of the gross (in one published structure, the storefront operator keeps 30 percent of gross and carries no payout risk; in another, roughly half of net) while absorbing payout liability at portfolio level — diversification across hundreds of firms smoothing the hot-cohort risk that kills single firms. None of these engines is the scandal. Each is a toll booth placed at a moment of maximum willingness to pay — the moment of failure, the moment of triumph, the moment of ambition — and a reader who can see the booths can price the

road.

Chapter 3 — Breakage: The Quiet Harvest

The challenge funnel takes its money at the door, in the open. The second great funding pattern of the rewards economy takes it at the exit, in silence. Its name in the accounting literature is *breakage*: the share of prepaid or earned value that is never redeemed.

Gift cards are the plainest case. Customers pay in advance; somewhere between 10 and 19 percent of gift-card value is never spent, by the Mercator figures carried through the industry literature. Loyalty points run higher: by the accounting literature's estimates, roughly 15 to 40 percent of point value expires or is abandoned unredeemed. During the holding period, the issuer has the use of the customer's money — an interest-free loan from the customer base. At expiry, under the revenue-recognition standards (ASC 606 and IFRS 15), issuers may recognize the expected unredeemed share as income, in some cases proportionally as redemptions occur — which is to say, the evaporation can be booked as revenue before it has even finished evaporating.

Notice who funds this system. No sponsor writes a check for the breakage line. The participant funds it, invisibly, by forgetting: by losing the card, by letting the points lapse, by moving house or changing phones or simply never checking a balance. The design does not require anyone to lose a contest. It requires only that a predictable fraction of people not collect what they already hold — and the fraction is predictable enough to be budgeted. A useful way to see the whole pattern: a conventional

breakage program is a promise issued at full face value, kept at a statistical discount, with the discount retained by the issuer as profit and printed nowhere a participant will read it.

The mechanics vary; the flow does not. Airline miles devalue by rule changes and expire by inactivity. Coffee points vanish at twelve months. Survey rewards expire a year after issue. In each case the value was earned or paid for, the expiry was disclosed somewhere in the terms, and the share that evaporates converts what was presented as a marketing cost — the price of loyalty — into issuer income. The law, for its part, mostly watches: federal rules guarantee *purchased* gift cards a five-year minimum life and restrict dormancy fees, but loyalty, award, and promotional value is expressly outside those protections, leaving the largest breakage pools governed by state unclaimed-property law and by whatever the issuer printed in its terms. Part I's first two systems took money from participants who were actively trying to win. Breakage takes it from participants who have already, in every ordinary sense, earned it. The participant funds this system too — and funds it most reliably at the moment they stop paying attention.

This matters for the design half of the book because breakage is not inherently dishonest. People will forget points in any system ever built; the forgetting is a fact of human nature, not a defect a designer creates. The design question is narrower and sharper: *when value evaporates, who receives it, and was that destination printed when the value was issued?* Part III's Reverse-Breakage Commons exists entirely inside that question. Hold it in mind through the next chapter.

Chapter 4 — Rake, Fees, and the Second Engine

Not every mechanism in the catalog is a trap. Part I ends with the mechanics that are *not* inherently unfair — because a reader who can only spot the bad ones has learned half the skill. Each of these has an exact point where it turns, and the point is always the same two: a rate that is not disclosed, or a discretion the operator holds over the outcome.

Tournament rake is the cleanest mechanic in the catalog. Entrants pay into a prize pool; the operator keeps a printed percentage and has no interest in who wins. Skillz, the skill-gaming platform, realized about 19.6 percent of gross volume as revenue in the third quarter of 2025 ($139.5 million of volume, $27.4 million of revenue, per its SEC filings); at a 25 percent rake, a head-to-head player must win about two matches in three just to break even. That is a hard price, but it is a *priced* price: the house's take does not depend on misleading anyone about it. The rake turns at exactly one point — when it is buried in prize-table arithmetic instead of printed as a number.

Platform commissions and fees work the same way when printed. Prediction exchanges charge per-trade fees by published formula — Kalshi's scales with price uncertainty, about $0.07 \times P \times (1–P)$ per contract — and creator platforms keep a published share of what flows through them: OnlyFans 20 percent, Patreon roughly 8 to 12 percent, YouTube about 45 percent of advertising revenue on the platform side. The fee formula is the disclosure; where the formula is public, the mechanic is a business. **Data and secondary revenue** — the second engine

whose toll booths Chapter 2 mapped — follows the same rule and breaks it more often. A firm may earn from selling back the very constraints it designed — split boosts, drawdown boosts, faster payouts — and, in FTMO's disclosed case, from trading real money on the data of its best simulated traders. None of these is unfair in itself. Each turns at the same point: when the second engine is invisible to the participant whose activity powers it, or when a rule sold as discipline is in fact a discretion — a termination clause the firm can invoke at will, a "gaming" definition applied after the fact. The line between risk management and rigging is whether the rule was published in advance and applied uniformly. Everything in Part II follows from that sentence.

The catalog also contains honest adaptations of the extractive mechanics, and they belong in this chapter rather than Part III because their participants *can* still lose money — a fact this book states rather than apologizes for. A season pass (pay per season, unlimited attempts within it) aligns the firm's revenue with tenure instead of failure count, and is fairer than per-attempt pricing. A retry priced far below the first entry is honest; a reset priced at 90 percent of a fresh purchase is the prop tell. A published rake of 10 percent or less on player-versus-player pools keeps the house indifferent to outcomes. A tier ladder — sustained performance unlocking bigger pools and better splits — costs a game operator almost nothing to grant, since simulated standing is free to create, and is among the most powerful retention designs in the catalog. Referral rewards are honest with a disclosed percentage and a cap; they turn extractive when they go multi-level, or when referrers profit specifically from their referees' losses. Strategy

royalties — a printed percentage of a copier's entry flowing to the player whose strategy was copied — turn skill into an asset for the skilled instead of a free input for the house, and are honest the moment the split is disclosed. A published reserve fund — prize pools held in disclosed, yield-bearing form, the yield split printed — borrows the insurance model's discipline without its opacity, and is honest exactly as long as the reserve's size, holdings, and split stay public. A fee refunded on success — FTMO refunds its two-step fee with the first payout — costs little precisely because few reach it (about 7 percent of one fee per buyer, in expectation, at Chapter 1's per-account payout rate) and buys an outsized fairness reputation; it is the single most transferable mechanic in the catalog. The catalog's cautionary middle deserves one entry of its own: the prize-arcade redemption economy, in which players pay per play and winners receive tickets redeemable for prizes at an exchange rate the operator sets. That model runs on three stacked spreads — play price, ticket yield, and prize cost — each invisible to the player, and claw machines add a fourth: a programmed payout rate, gripping hard only often enough to hit a target. It is honest only if the ticket-to-prize exchange rate and prize values are posted, which is to say it is a working model of everything the designs in Part III invert. These are systems a participant can lose in, entered with open eyes and printed numbers. They are not zero-loss, and Part III is reserved for designs that are. One design, however, earns this chapter's capstone, because it does something none of the others do: it repairs the challenge model itself.

The Completion Underwriting Market

[ORIGINAL DESIGN] — presented here, in Part I, for a reason this section will state in full.

The prop model's core conflict is that the house sets the odds and holds the other side: the firm profits when its customers fail, and the completion rate — the single number that prices everything — is the industry's best-kept secret. The Completion Underwriting Market removes the conflict by publishing the number and selling the risk. A challenge series runs on refundable rules: entry is free, or the entry stake is a refundable lock of the kind Part III's Mastery Bond describes, so no player's money is consumed by entering. Each completer earns a fixed bonus, published in advance — say $25. A sponsor who wants those bonuses guaranteed pays a premium, priced from the series' **published historical completion rate**. Independent underwriters — not the operator — lock collateral at a printed multiple of the expected bonus cost and hold the variance. The operator takes a flat series fee, identical whether nobody completes or everybody does, and is **barred from underwriting its own series**.

The worked arithmetic: 1,000 entrants, a $25 bonus, a published historical completion rate of 12 percent. Expected bonus cost is $3,000; the premium is set at 110 percent of expected, $3,300; underwriters lock $6,000, twice the expected cost. If the cohort completes at 12 percent, bonuses are $3,000 and underwriters keep $300. If it runs cold at 8 percent, bonuses are $2,000 and underwriters keep $1,300. If it runs hot at 20 percent — the gold-rally scenario of Chapter 2 — bonuses are $5,000 and underwriters lose $1,700 of stake. They eat the hot cohort exactly as FTMO ate its 2025 cohort, except that they volunteered, at a published price, against published odds,

holding no control over the gate. The sponsor's total cost is $3,550, known in advance. The operator's interest in any player's failure is zero.

Now the audit this design received before it was allowed into this book, stated as plainly as the design itself. Tested seat by seat against this book's definition — *no participant can end with less money than they started with* — the player seat qualifies fully: entry is free or refundable, and the bonus is purely additive. The sponsor seat qualifies as a payer: its cost is a fixed price known in advance, and it cannot lose more than the printed price. **The underwriter seat does not qualify.** An underwriter stakes capital that can be lost — $1,700 in the worked hot cohort — and in a hot cohort the bonuses above the premium are, in fact, paid out of underwriter losses. The design's defense is a real one, and it is also a limited one: underwriters are professional risk counterparties, the same seat insurers and sureties occupy throughout the prize industry, paid a premium to hold priced, published variance, choosing their exposure series by series. That defense holds only under a printed condition, which is part of the design and not a footnote: **underwriting must remain a professional seat** — priced, collateralized, opt-in per series, worst case printed. If underwriting is ever opened to retail stakers as a yield product, those stakers become participants in substance, and the design fails the definition in fact.

A second audit finding travels with the design wherever it is used: **the guarantee is bounded.** Bonuses are covered only up to the premium plus the locked stake. An honest series must print its rule for completions beyond that bound before entry — the audit's recommendation is a printed series cap, with

completions beyond the cap sharing the remaining fund pro-rata. A guaranteed bonus with an unprinted ceiling would fail this book's own standard. The feasibility record is equally part of the design: the market has a cold-start problem (no underwriter prices a series without audited rate history, and the operator is barred from seeding its own market); adverse selection runs both ways (underwriters will prefer predictable pools and shun leaky ones — which is also the design's disciplining virtue, since an uninsurable gate is information); completion risk inside a cohort is correlated, not independent, because one leaked strategy can move a whole cohort hot at once; and the largest barrier is regulatory — insurance and surety are licensed activities, and an open pool of staking underwriters needs licensed structure or counsel's alternative before real money moves.

That is why this design closes Part I instead of opening Part III. Part III stays zero-loss-pure. This design keeps priced risk in the system — held by professionals, in the open, at a published price — and in exchange it removes the one feature that makes the challenge industry what it is: a house that profits from failure and prices its secret in the dark.

Closing section — Price the Gate

Strip the design to its principle and a simpler practice falls out, available to any operator today: **publish the completion rate and let the market price the difficulty.** Every challenge series already has a price for its gate — the fee, the reset schedule, the add-on menu — but the price is set against a number only the house can see. A series that published its historical completion rate would find its gate priced by outsiders within a season: sponsors deciding what a guaranteed bonus is worth,

underwriters deciding what the variance costs, players deciding what an attempt is worth to them. The industry's most valuable secret is a single percentage. Everything in Part I has been a study of what that secrecy buys. Part II begins with what publishing buys instead.

Part II — Principles of No-Loss Design



Plate — Your Named Star

Art: Douglas John Brownfield

Chapter 5 — Who Pays, and Why Paying Is Worth It

Zero-loss does not mean money from nowhere. It means the participant is never the source of the prize. Across the twenty models in Part III, every payer reduces to one of ten archetypes, and the archetype matters because it tells the designer where the money's loyalty lies — what the payer is actually buying, and therefore what would make them stop.

1. **Merchants and brands** pay because a delivered, verified customer is worth more than the reward that delivered them. (Models 1, 11.)
2. **Endowers and foundations** pay because a name and a mission funded forever, from returns alone, is worth a one-time gift of capital. (Model 2.)
3. **Advertisers, through content revenue,** pay because attention is inventory, and sharing its revenue buys an audience's loyalty. (Model 3.)
4. **Patrons** pay because the board or the ladder is itself their marketing and their meaning — visible, named, narratable generosity. (Models 4, 8.)
5. **Peers, in kind,** pay each other: an hour given is an hour receivable, and the ledger is the product. (Model 5.)
6. **Data buyers** pay because aggregated, consented data answers questions worth money. (Models 6, 17.)
7. **A sister business** pays a printed percentage of a working margin because the games build the community that becomes its customers. (Model 7.)

8. **Time itself** pays, through yield: pooled money earns while it waits, and the float funds the prize. (Model 9.)
9. **A marketplace's users,** through commissions, pay because skilled activity is what makes the marketplace's tiers and goods valuable. (Model 10.)
10. **Cost-avoiders and outcome funders** pay because a verified avoided cost, or a delivered outcome, is worth a printed share of what it saved or accomplished. (Models 12, 13, 14.)

Around the ten stand four more payers the models use in supporting roles: employers, who fund learning because trained people are worth more than tuition (Model 18); cooperative members in their capacity as owners, returning their own surplus by formula (Model 16); venues holding perishable inventory, for whom a filled seat beats an empty one at any price above zero (Model 15); and buyers of idle capacity, who pay for computing and bandwidth that already exist (Model 19). Model 20's matching-pool funders are endower and patron money, allocated by community signal rather than by the funder's own choice.

Payer archetype	Why paying is worth it	Models
Merchants & brands	A delivered, verified customer is worth more than the reward	1, 11
Endowers & foundations	Perpetuity: a name and a mission funded forever from returns	2
Advertisers (via content revenue)	Attention is inventory; sharing its revenue buys loyalty	3

Payer archetype	Why paying is worth it	Models
Patrons	The board or ladder is the patron's marketing and meaning	4, 8
Peers (in kind)	An hour given is an hour receivable; the ledger is the product	5
Data buyers	Aggregated, consented data answers questions worth money	6, 17
A sister business	A printed percentage of a working margin buys community and goodwill	7
Time itself (yield)	Pooled money earns while it waits	9
A marketplace's users (commissions)	Skilled activity makes the marketplace's goods valuable	10
Cost-avoiders & outcome funders	A verified avoided cost or delivered outcome is worth a printed share	12, 13, 14

One exclusion belongs in this chapter, stated out loud rather than left implicit: **the participant's own future earnings are not a payer in this book.** Income-share designs — fund the player's entry or training now, collect a percentage of what they earn later — were considered and excluded, because a bad outcome can leave a participant owing more than they received, and a design that can end there fails the definition no matter how gently the collection is framed. Governments, likewise, appear in these models as outcome buyers and utility payers, not as direct prize payers; the archetypes above are the ones the

evidence supports.

The chapter's rule is short enough to print on the inside of a door: **a no-loss design is a payer with a real reason, plus arithmetic that survives a bad month.** The payer's reason is the foundation — if paying is not worth it to the payer on an ordinary month, the design is a subsidy with a marketing department, and it will end the way subsidies end. The bad-month clause is the discipline: every model in Part III is shown at a small scale with its fragile point named, because a funding structure that only works in its best month is the fee pool wearing better clothes. The archetypes differ in how their reasons weather a bad month, and the difference is worth reading before choosing one. An endowment's reason is structural — the capital is already given, and the spending rule shrinks the prize rather than breaking a promise. A merchant's reason is transactional and renews with every pool the merchants fund: it survives exactly as long as the visits keep verifying. An advertiser's reason is cyclical, and a design built on attention revenue must print that the dividend varies with revenue, because it will. A cost-avoider's reason is durable while the inefficiency lasts — the saving is real money the payer was about to spend — and it expires when the inefficiency does. None of this ranks the archetypes; it tells the designer which question to ask of each. When you evaluate any system — one in this book or one offered to you — start here. Name the payer. State their reason in one sentence a skeptic would accept. Then ask what happens to the prizes in the payer's worst month. If the answers hold, the design deserves the next chapter's test. If they do not, no amount of polish further down will save it.

Chapter 6 — Transparency Is the Product

A fair funding structure can still be run dishonestly, and the difference lives in the pipeline: who holds the prizes, who watches them move, and what is printed where participants can read it before they commit. This chapter states the standard the models in Part III are held to: printed rates, published pools, verifiable draws, and independent control of the prize pipeline.

Printed rates means the numbers that govern a participant's outcome — the fee, the share, the split, the expiry rule, the completion rate where one exists — are published before entry, in the same document that invites entry, at a size a reader will actually encounter. Published pools means the money's state is observable: a pool's balance, its funding source, and its movements are reported on a schedule participants can check, not summarized in an annual boast. Verifiable draws means any random selection can be re-performed or audited by someone other than the operator. Independent control means the hands that award prizes are not the only hands that can reach them.

The cautionary case is McDonald's Monopoly, and it is the right case precisely because its funding was clean. The promotion, running since 1987, was a genuinely zero-loss game: prizes were funded from the marketing budgets of the restaurant business and its supplier partners, and a free mail-in entry route existed, so no customer had to buy anything to play. By this book's definition, the design was fair. What failed was the pipeline. From 1989 to 2001, Jerome Jacobson, the security director of the promotion's administrator — the vendor charged

with guarding the high-value game pieces — stole those pieces and distributed them through a network of intermediaries. More than $24 million in prizes was diverted. More than 50 people were convicted. McDonald's paid a $10 million make-good giveaway and $16.6 million to settle with its administrator. For over a decade, the game's funding was honest, its advertising was accurate, its prizes were real — and its winners were chosen in a parking lot by the man hired to prevent exactly that.

The lesson generalizes beyond one promotion. Fair funding with an unaudited pipeline is not fair; it is merely fair-*funded*, and participants experience the two as identical until the audit that never happens. This is why Part III's designs keep returning to the same furniture — published pool balances, printed formulas executed as printed, verifiers separate from operators, reserves reported movement by movement. The standard is concrete enough to check against any running program. A board that posts its bounty and its resolution source in advance (Model 8) can be audited by anyone who can read. A commons that publishes its pool balance and its sweep rule at issuance (Model 11) leaves its participants nothing to take on trust. A float pool that names its custodian and its current yield before entry (Model 9) shows the reader four numbers that are the entire business model. In each case the test is the same: could a participant, using only what the program publishes, re-derive what they were promised and what they received? Where the answer is yes, transparency is doing its work. Where the answer requires the operator's private ledger, the program is asking for trust it has not earned — whatever its funding structure claims. Transparency is sometimes described as a trust feature, a nice signal a good operator sends. That understates it. In a zero-loss

design, transparency is load-bearing: it is the mechanism by which a participant verifies that the payer is who the label says, that the share is what the label promises, and that the prize they were shown is the prize that will be awarded. It is, in the plainest sense, the product.

Chapter 7 — The Consideration Point

Zero-loss designs are also, not coincidentally, the legally simpler designs — and the reason is a single word in gambling law: *consideration*.

Under United States state law, an illegal lottery generally requires three elements together: a **prize**, **chance**, and **consideration** — something of value given to enter. Remove any one element and the lottery analysis fails. Every model in Part III removes consideration from the participant side: nobody pays to enter. That is not a legal trick layered onto the designs after the fact; it is the same fact as the definition, seen from the law's side of the table. A system in which participants cannot lose money is, structurally, a system in which the gambling question mostly does not arise.

"Mostly" is doing work in that sentence, and this chapter keeps it. What follows is headline-level orientation for United States programs — the legal treatment in this book is US-only by design — and it is a map of where a lawyer's look is worth paying for, not legal advice.

Where **chance** decides winners, a free-entry promotion is a sweepstakes and carries its own duties: publish official rules and odds; provide a free alternative method of entry with equal dignity wherever any other route involves a purchase; and mind the registration thresholds that arrive quietly as a program grows — New York and Florida require registration and a bond or trust for prize pools over $5,000 (New York filings are due at least 30

days before launch; Florida's at least 7), and Rhode Island requires registration for retail promotions over $500. Skill-based designs sit differently, because attainment removes chance rather than consideration; the models in this book use both routes deliberately. One open question deserves a flag rather than an answer: a few states scrutinize whether substantial effort, or personal data itself, can count as consideration. That question is unsettled and state-specific, and any program that runs on heavy participant effort should have it checked for the states it actually operates in.

Stored value carries its own exposure. Every state has unclaimed-property law: dormant balances are generally reportable to the state after a dormancy period, commonly three to five years. Federal gift-card rules set a five-year minimum expiry and restrict dormancy fees for *purchased* gift cards — but loyalty, award, and promotional cards are excluded from those federal protections, which means most programs in this book operate under state law alone. Any model that holds participant value over time — the float pool, the mastery bond, the breakage commons above all — needs a per-state review, and the Commons' dormancy sweep is lawful only if the swept value is not the owner's property under the applicable state's law. That is the single most important legal question in this book's model set, and Chapter 15 treats it that way.

Securities law asks its own question, the Howey question: is there an investment of money in a common enterprise with an expectation of profits from the efforts of others? Most models in this book avoid it because participants put no money in. The exposed points are named where they occur: pooled participant money earning yield (Model 9, mitigated by its principal-back

form and by treating yield as prizes rather than promised returns), and any adaptation that lets participants buy into a pool or pays them in a transferable token rather than cash, credit, or prizes. An endowment is safe on this point precisely because participants never contribute to it.

Tax reporting is the most concrete section, and the figures here were verified against current IRS publications. Prizes and awards are taxable income to the recipient. Payers generally report them on Form 1099-MISC. The reporting threshold was $600 for decades; under §70433 of the One Big Beautiful Bill Act (P.L. 119-21), it is **$2,000 per payee per calendar year for payments made after December 31, 2025**, inflation-adjusted for later years — the same $2,000 threshold covers nonemployee compensation on Form 1099-NEC, including prizes for services. Two carve-outs did not move: royalties remain reportable at $10, and gross proceeds paid to an attorney at $600. And a payment below the reporting threshold is still taxable income to the recipient — only the payer's filing duty changes. Patronage dividends are reported on Form 1099-PATR at $10 and above; employer education benefits are tax-free to the employee up to $5,250 a year under IRC §127, with a written plan. Barter income, finally, is taxable at fair market value in the year received — the Internal Revenue Service's general published position, stated here without the stronger exemption claim that circulates in time-banking literature and does not survive the primary sources; Chapter 14 gives the narrow, accurate version.

Privacy and data duties attach wherever a model sells or uses participant data and behavior — principally Models 6, 13, and 19. Consent and opt-out duties under California's CCPA/CPRA

"sale/share" framework, and lawful-basis and purpose-limitation requirements under the GDPR for any European participants, set the floor; health-adjacent data, Model 13's territory, raises the stakes further. Consent that names the dataset and the buyer class is both the ethical requirement of this book's designs and the direction of the law. **Money handling** is the last flag: models that hold or redistribute cash for others — the float pool, the mastery bond, the bounty board's poster side — can implicate state money-transmission licensing depending on structure. Using established payment processors and avoiding stored cash balances is the usual small-scale mitigation, and it is a counsel question, not a design assumption.

The chapter's point is not that these designs are law-free. It is that the legal work concentrates in known places — sweepstakes duties, stored value, the seals around pooled money — and that a design which begins by removing the participant's money from the table has already answered the hardest question the law asks of a game.

Chapter 8 — Name the Contribution

Zero-loss is a claim about money. It is not a claim about costlessness, and a book that let the two blur would be doing the thing it criticizes. Participants in every model in Part III contribute something real. The principle of this chapter is that the contribution belongs on the label — named, in plain words, at the point of entry — because a contribution a participant cannot see is a price they did not agree to pay.

The contributions in this book's models, in ascending order of weight. **Attention**: watching, reading, searching, returning — worth cents per participant-hour at real advertising rates, and honest only when the design says the amount varies with revenue. **Time**: hours of study, play, teaching, and participation; the Patron's Ladder and the Training Purse run on it, and the Skill Barter Bank runs on nothing else, at a printed par of one hour for one hour. **Judgment**: reading projects and signaling support, the whole of the Quadratic Commons' ask. **Presence**: showing up — foot traffic for a merchant, a filled seat for a venue, the audience whose existence justifies an endowment's season. **Behavior**: changing what you do at a measured moment — using less electricity at the expensive hour — where the measurement itself is also a contribution. **Data**: purchase histories, study responses, telemetry — specified, consented, and revocable, or it is not a contribution but an extraction. **Work**: producing a verified outcome in the world. **Device capacity**: computing power, bandwidth, storage — plus the electricity that runs them and the wear they take. **Purchases at**

fair prices: the Patronage Pool's honest core — the member is a customer, the reward is their own surplus returning, and the design must not pretend otherwise. **Temporary lockup**: money at rest in a refundable stake — recoverable in full, and still a real cost while it rests, which is why the Mastery Bond names it rather than waving it away.

Two models in this set carry heavy contributions, and this book says so without flinching, here and in their chapters. The Non-Event Dividend (Model 13) runs on continuous verified behavior — the most surveillance-heavy contribution in the set: a participant trades ongoing proof of how they live for rewards funded by their own risk reduction. The Idle Engine (Model 19) draws on a participant's devices, electricity, and bandwidth — the most physically real contribution in the set — and its economics can quietly go negative for the participant if buyer rates fall or power draw rises, which is why its metering must show the participant their own arithmetic and not just the gross. Neither model is disqualified by its weight. Both are disqualified the moment their weight is hidden.

The test is portable, and it is the same test as Chapter 5's, from the other chair. Before entering any system, a participant should be able to complete the sentence: *what I give is ___; what it earns is ___; the share was printed at ___*. A system that can fill in all three blanks, in words its participants actually use, has named its contribution. A system that cannot has told you what it is.

The naming has mechanics, and the models show them working. The Research Bounty Board names the contribution *before* commitment — pay and time estimate on the listing, the posted

reward equal to the amount received — so consent is informed at the only moment it can be. The Data Dividend makes the specification itself the consent: what is collected, what is sold, to whom, in what aggregation, revocable when the participant says so. The Idle Engine's conditions — idle, charging, unmetered — are printed as part of the bargain, and its dashboard reports the participant's own side of the arithmetic. The Mastery Bond names lockup as a cost even though the money returns in full, because recoverable is not the same as free. In each case the mechanism is identical: the contribution appears in the same document, at the same size, as the reward it earns. That placement — not the existence of a disclosure somewhere in a terms file — is what "named" means in this book, and it is the standard the Conclusion's checklist will ask about by name.

Chapter 9 — For the Builders, and the Players

Everything so far converges on one standard, and it can be read from two chairs. This chapter reads it from both, directly, before the designs begin — because the twenty models that follow are twenty ways of meeting it, and a reader who holds the standard can judge each model, and any system outside this book, without the author's help.

To the builder. You owe the people in your system three things, and none of them is optional or decorative. First, a *printed flow*: who pays, at what rate, posted where entrants can read it before they commit — not in a filing, not in a help page three clicks deep, but in the document that invites them in. If your funding cannot survive being printed on your own front door, the funding is the problem, not the printing. Second, a *named contribution*: whatever your participants give — time, attention, data, behavior, lockup — stated on the label, in those words, with its weight honestly described. The sentence "free to play" is a claim about money; do not let it smuggle a claim about everything else. Third, an *inspectable prize pipeline*: reserves and payout paths a participant can examine, under control independent enough that the failure in Chapter 6 — a fair game stolen for a decade by the man guarding its prizes — cannot repeat quietly inside yours. A system that meets all three can charge its payers proudly and publish its rates like a menu. A system that meets two has a named defect; fix it or print it. A system that meets none is a funnel with good graphic design.

To the player. Before you give any program your money — and before you give it your time, attention, or data, which are also yours — demand the same three things from the other chair. *Show me the flow*: who pays for the prizes, and why is paying worth it to them? If the answer is a fog of words like "ecosystem" and "rewards pool" with no name and no reason attached, the answer is you. *Name my contribution*: what exactly do I give, and what is it worth to the people receiving it? A program that cannot say what your attention or your data earns it is not being modest; it is keeping the ledger on one side. *Let me inspect the pipeline*: where is the money held, who can reach it, and who besides the operator has checked? A prize awarded by an uninspectable process is a story, and stories are cheap to produce at scale. These demands have documents attached to them, and asking for the documents is the practical form of the demand: the fee table and the official rules, for the flow; the program's own statement of what participants contribute and what it earns the operator, for the contribution; the pool ledger, the reserve report, the custodian's name, for the pipeline. A fair program produces these the way a restaurant produces a menu — as the ordinary furniture of doing business in the open. A program that treats the request itself as unreasonable has graded its own paperwork.

The two halves are one standard. The builder's duties are the player's demands, word for word, and a fair system is simply one whose paperwork survives being read from both chairs at once. If a system's documents cannot survive that reading, it is not fair play, whatever it calls itself — and the finding does not depend on anyone's intentions, which is the point of having a standard instead of a vibe.

The Conclusion of this book compresses this chapter into a five-minute checklist — six questions, applicable to any program in the time it takes to drink a coffee. This chapter is the address; the checklist is the tool. Read the models that follow with both in hand: the builder's question for each is *could I run it to this standard*, and the player's is *would I enter it under this standard*. The designs that satisfy both readers at once are the ones worth building.

Part III — The Models: Twenty Worked Designs



Plate — Every Port of Call

Art: Douglas John Brownfield

Chapter 10 — Funded by Marketing

Three models in which the payer is buying customers — the most familiar honest reason in this book.

Model 1 — The Merchant's Table [FOUND]

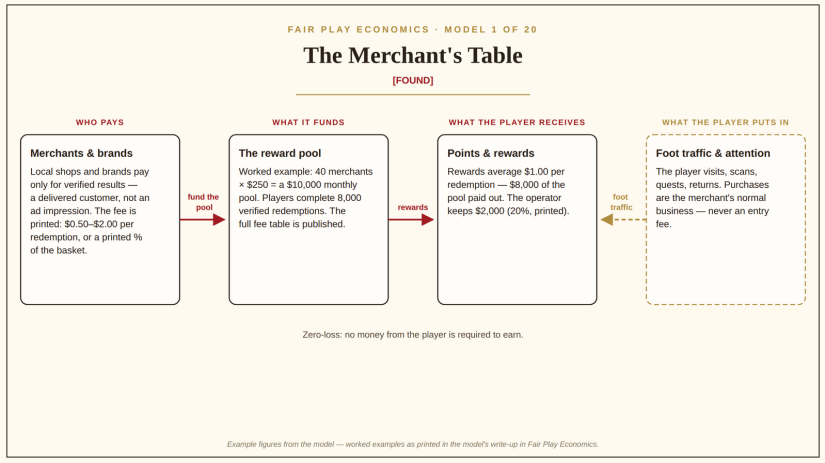


Diagram — The Merchant's Table: who pays, what it funds, what the player receives, what the player puts in.

The payer and their reason. Local merchants and brands. A customer who walks in because of the game and redeems an offer is a delivered customer — the same result merchants already pay coupon networks, card-linked-offer platforms, and delivery apps to produce, except that here they pay per verified result, not per impression.

The flow. A publisher runs quests — visit, scan, buy anything, return visit. Completing a quest earns points or direct rewards from a pool the merchants fund. The merchant pays the operator

a printed per-redemption fee (say $0.50 to $2.00) or a printed percentage of the redemption basket, and the operator funds rewards from those fees at a printed rate. The full fee table is published.

The arithmetic. A café group funds a monthly pool: 40 merchants × $250 = $10,000. Players complete 8,000 verified redemptions; rewards average $1.00 per redemption ($8,000), and the operator keeps $2,000 — 20 percent, printed. A merchant whose $250 bought about 200 redemptions paid $1.25 per delivered visit, cheap against typical local advertising, with every cent accounted to a visit that happened.

The contribution, named. Foot traffic and attention: the player goes somewhere and does something. No money from the player is required to earn; purchases at the merchant are the merchant's normal business, not an entry fee.

The precedent. Partial analogue. Fetch, founded in 2013, runs the brand-funded half at national scale: brands pay commissions and fund bonus points when shoppers scan receipts showing their products, and pay again for the aggregated purchase data — a General Mills executive described receiving about 62 million lines of data per day from the partnership. Fetch proves merchants and brands will fund consumer rewards at scale when redemption is verified. The quest form itself — rewards for verified visits and tasks, funded per redemption — is also running in the world: Niantic's sponsored locations in Pokémon GO charge sponsors on a cost-per-visit basis, and shared merchant pools paid out by a published rule exist in cashback platforms (GenieGenerate runs one, filled by local shops' cashback and paid by a printed rule). The mechanic is found, in

use, on both halves.

The verdict. Workable now. Verification tooling exists off the shelf — Fetch is the existence proof for the verification half. The merchant side clears on ordinary numbers: $1.25 per delivered visit, all-in, against typical local baskets of $8 to $15, works if even a third of quest visits are incremental. The model's fragile point is merchant churn if attribution is disputed — publish the verification method alongside the fee table. Quests decided by chance need the sweepstakes treatment of Chapter 7, and paying per purchase must never make a purchase the only way to earn: keep a free quest path.

Working exhibit: The Café (freedough.pages.dev/#arcade), aboard FreeDough — the café and shops fund the games because the games bring the crowd.

Model 8 — Sponsor-the-Board [FOUND mechanic / ORIGINAL fee-wrapper packaging]

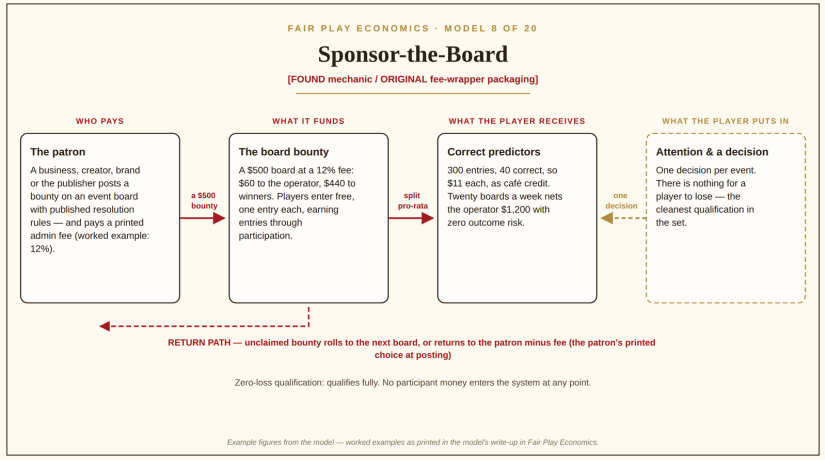


Diagram — Sponsor-the-Board: who pays, what it funds, what the player receives, what the player puts in.

The payer and their reason. A patron — a business, a creator, a brand, or the publisher — posts a bounty on an event board with published resolution rules, and pays because verified attention, and, where winnings are paid as store credit, verified foot traffic, is worth a known price. The operator charges the patron a printed admin fee.

The flow. Players enter free, one entry each, earning entries through participation. Correct predictors split the bounty pro-rata. An unclaimed bounty rolls to the next board publicly, or returns to the patron minus the fee — the patron's printed choice, made at posting.

The arithmetic. A $500 board at a 12 percent fee: the operator takes $60, and $440 reaches the winners. With 300 entries and 40 correct predictors, each receives $11 as café credit. Twenty boards a week nets the operator $1,200 a week with zero outcome risk.

The contribution, named. Attention and one decision per event. There is nothing for a player to lose; this is the cleanest qualification in the set, since no participant money enters the system at any point.

The precedent. None found as a standing system; partial analogues exist. Brand sweepstakes and radio promotions run one-off versions of sponsor-paid prizes; prediction markets run the board without the sponsor; free-entry, sponsor-funded prediction play exists in sweeps formats. The standing, rules-published, rollover-accounted board — with the patron paying a printed fee for the posting itself — is the element that appears unused, plausibly because patron acquisition is a sales job and prediction economics chase trading volume instead.

The verdict. Workable now. The constraint is patron sales, not mechanics — at $60 a board in admin fees, resolution and anti-bot operations are covered at modest volume, and the operator carries zero outcome risk throughout. Where winners among correct predictors are chosen by a tie-break draw, chance enters and sweepstakes duties apply, including the registration thresholds of Chapter 7 as boards grow. Two operating rules carry the design: the resolution source must be named and contracted in advance, and free entries invite automation, so entries should be tied to participation that is costly to fake.

Working exhibit: The Board (freedough.pages.dev/board), aboard FreeDough — sponsors post bounties, and players claim them by answering.

Model 15 — The Empty Seat Exchange [ORIGINAL DESIGN; partial analogues]

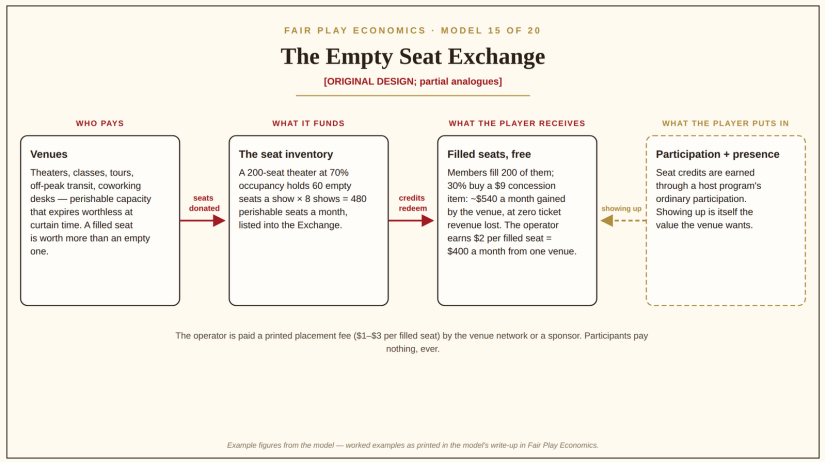


Diagram — The Empty Seat Exchange: who pays, what it funds, what the player receives, what the player puts in.

The payer and their reason. Venues and operators of perishable capacity — theaters, classes, tours, off-peak transit, coworking desks — hold inventory that expires worthless at curtain time. They donate seats as rewards because a filled seat is worth more than an empty one: concessions, audience energy, reviews, future sales. The operator is paid a printed per-redemption placement fee by the venue network or a sponsor.

The flow. Participants earn seat credits through a host program's ordinary participation — quests, learning rungs, community tasks. Venues list perishable inventory into the Exchange; credits redeem first-come. The venue pays nothing to list; where no sponsor carries the fee, the venue can pay in kind — a share of listed inventory for the operator's own community.

The arithmetic. A 200-seat theater runs at 70 percent occupancy: 60 empty seats a show, 8 shows a month, 480 perishable seats. If Exchange members fill 200 of them and 30 percent buy a $9 concession item, the venue gains about $540 a month in concessions plus 200 audience members it did not have, at zero ticket revenue lost — the seats were empty. The operator, at $2 per redemption, earns $400 a month from one venue. Participants pay nothing, ever.

The contribution, named. Participation in the host program, and presence — showing up is itself the value the venue wants.

The precedent. Partial analogue, with the funding direction inverted. Fill-a-seat membership clubs distribute unsold tickets to members, but members pay subscriptions for access, which makes the member the payer. Broadway-style digital lotteries allocate cheap or free seats by chance as marketing one-offs.

The Exchange — seats as *earned rewards* inside an unrelated program, with the operator paid per redemption — was not found operating.

The verdict. Workable now. Perishable-inventory donation is the cheapest payer contribution in the entire set, since the inventory's alternative value is zero, which makes venue recruitment the easiest sales conversation in this book. The fragile points are prime-inventory leakage — venues must list genuinely perishable stock, released close to showtime, or members learn to wait for free seats instead of buying — and no-shows: a redeemed-but-empty seat is worse than an unlisted one, and refundable-credit no-show rules, borrowed from Model 10's bond logic, are the standard mitigation.

Chapter 11 — Funded by Time

Two models in which the payer is patience: capital given once and never spent, and money that earns while it waits.

Model 2 — The Endowment Season [ORIGINAL game-season arrangement; funding structure FOUND]

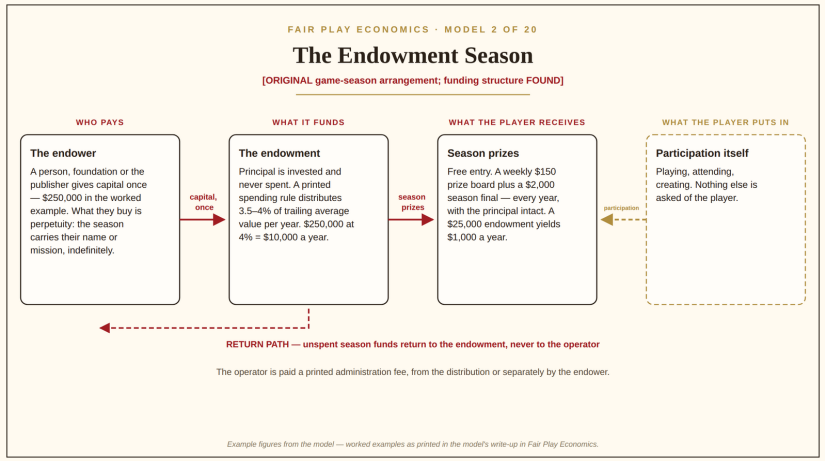


Diagram — The Endowment Season: who pays, what it funds, what the player receives, what the player puts in.

The payer and their reason. An endower — a person, a foundation, or the publisher itself — gives capital once. The endowment's investment returns fund a recurring season of prizes, indefinitely; the principal is never spent. The payer buys perpetuity: the season carries the endower's name or mission at a one-time cost.

The flow. The endowment is invested conservatively under a printed spending rule — distribute up to about 4 percent of trailing average value per year. Each season's prize fund is the spending-rule amount. Players enter free. Unspent season funds

return to the endowment, compounding future seasons, never to the operator; the operator is paid a printed administration fee from the distribution or separately by the endower.

The arithmetic. A $250,000 endowment at a 4 percent spending rule yields $10,000 a year — a weekly $150 prize board plus a $2,000 season final, every year, indefinitely, with the principal intact. A $25,000 endowment yields $1,000 a year: a modest but genuinely perpetual local season. The ratio is the model's honesty and its limit in one line: at a 4 percent rule, every $1,000 of annual prizes needs $25,000 of endowment, and no marketing changes that.

The contribution, named. Participation itself — playing, attending, creating. Nothing else.

The precedent. Full precedent for the funding structure, and the strongest feasibility evidence any model in this book has. The Nobel Foundation has run on this structure for 125 years: Alfred Nobel's bequest of about SEK 31 million in 1895, invested; invested capital stood at SEK 6,744 million at the end of 2025, and the 2025 prizes were SEK 11 million each. The Foundation's stated goal is a return of at least 3 percent above inflation to preserve the base, and the prize amount is set annually based on its finances — it cut the prize by 20 percent in 2012 when capital required it. A widely repeated claim that the Foundation follows an iron-clad rule of never withdrawing more than 3.5 percent could not be verified against the Foundation's own sources and appears only in social-media retellings; this book prints the target-and-annual-setting version, because that is the version the record supports. What appears original is the arrangement: a recurring *game prize season* funded solely by

yield on preserved principal. The funding structure is a century and a quarter old; the game-season application is this book's.

The verdict. Workable now; prize size is scale-bound. The structure is proven — the only variable is how much endowment a given publisher can raise. Governance follows endowment norms (in the US, UPMIFA-style spending rules for institutional funds). One boundary matters: participants are never invited to contribute to the endowment expecting returns; if they were, the securities analysis would change completely. In this design, they are not.

Model 9 — The Zero-Rake Float Pool [FOUND core concept / ORIGINAL flipped allocation — novelty unverified]

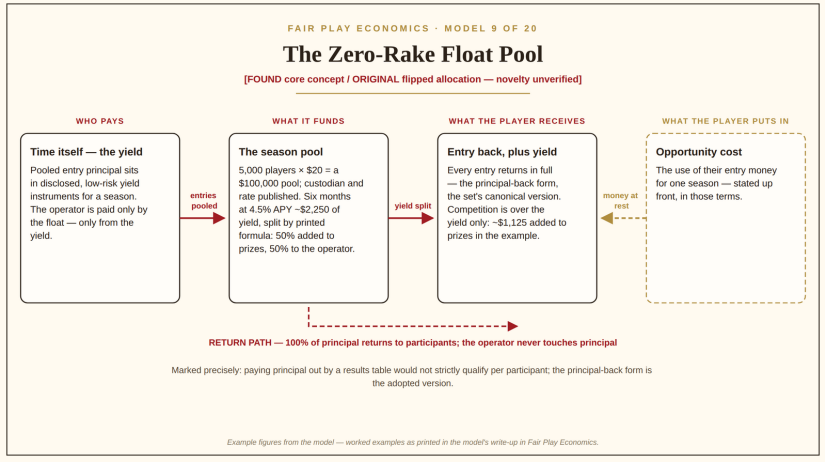


Diagram — The Zero-Rake Float Pool: who pays, what it funds, what the player receives, what the player puts in.

The payer and their reason. Time. Players' pooled entry principal sits in disclosed, low-risk yield instruments for a season; the yield is split by a printed formula. The operator is

paid only by the float — never by the principal, and never by a rake on prizes.

The flow. Entries are pooled with a published custodian at a published rate. In the form this book adopts — and the qualification matters, so it is stated here rather than footnoted — **every participant's entry is returned in full**, and competition is over the yield (and any sponsor top-up) only: at season end, principal goes back to its owners, and the yield is distributed as prizes under a printed split, part to prizes and part to the operator. Before entry, the player sees four numbers: pool size, custodian, current yield, split.

A second form of the design exists and does *not* strictly qualify: distributing the full principal as prizes by a results table, under which a player who places low receives less than their entry. That form's guarantee is pool-level — all principal returns to participants as a class, and the operator never touches it — but it is not a per-participant guarantee, and this book's definition is per-participant. The set therefore adopts the principal-back form as canonical and records the other rather than smoothing it over.

The arithmetic, stated honestly. 5,000 players × $20 = a $100,000 pool; a six-month season at 4.5 percent annual yield produces about $2,250; at a 50/50 split, the operator's share is about $1,125 a season. At 100,000 players — a $2 million pool — the operator's share is about $22,500 a season. At 0.5 percent yield it earns nearly nothing, and the design must say so on its face. Float pays for patience and size, and at small size it pays almost nothing; any small operator presenting it as a business is overselling.

The contribution, named. The opportunity cost of the entry money for one season, stated up front in those terms.

The precedent. Partial analogue. Insurance float — premiums invested before claims come due — is the funding logic at industrial scale, and the endowment of Model 2 is its cousin; refundable entries pooled for yield have live relatives in prize-linked savings products. The core concept — pooled refundable entries, invested, with the return as the prize component and the operator paid from the return — is found, including in patent territory dating to 2006. The possibly original element is a single allocation flip: principal returned in full, prizes drawn from yield, and the operator paid *only* a printed share of yield. That flip sits over a dense family of prior designs, and its novelty is unverified; this book claims the arrangement, not the invention. The design's natural home follows from its math: a large, patient, distributed community, or a publisher whose real revenue is earned elsewhere — books, content, a marketplace — and who wants a prize system that is provably not eating the players.

The verdict. Fragile at small scale; workable at scale as an ecosystem feature. The math is the verdict: float pays for patience and size, and at small size it pays almost nothing — any small operator presenting it as a business is overselling. The design also inherits a reputation problem it did not earn: yield-bearing customer pools were the shape of the opaque lending products that collapsed in 2022–23, and the honest version, disclosed custodian and all, will be read in their light until its own record says otherwise. Legal flags are the heaviest in the set: pooling participant money that earns yield invites securities analysis (the Howey question of Chapter 7) and

money-transmission and custody questions. The principal-back form, with yield treated as a prize pool rather than promised returns, mitigates the analysis but does not eliminate the need for counsel.

Chapter 12 — Funded by Attention and Data

Four models in which the participant's contribution is the product being sold — attention, data, study responses, and idle capacity — and in which naming that contribution is therefore not a courtesy but the entire design.

Model 3 — The Attention Dividend [FOUND]

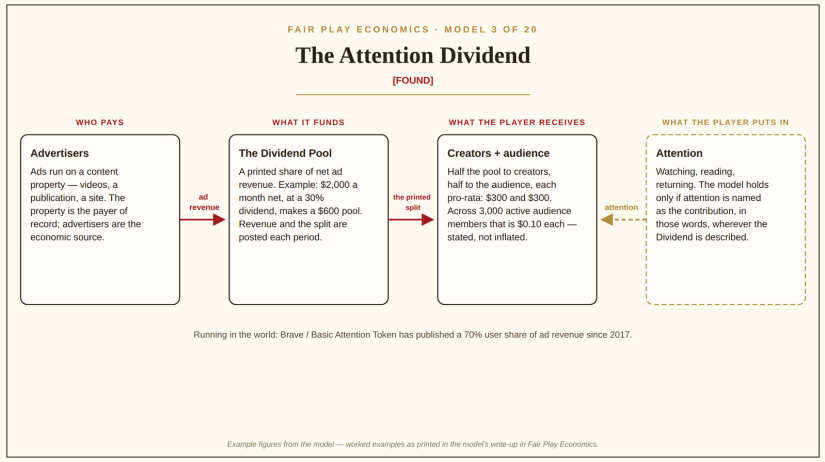


Diagram — The Attention Dividend: who pays, what it funds, what the player receives, what the player puts in.

The payer and their reason. Advertisers, through a content property's advertising revenue. The property — videos, a publication, a site — earns advertising money because an audience shows up; the Dividend prints the share of that revenue returned to the audience and the creators who drew it. The payer of record is the property; the economic source is the advertiser.

The flow. The property publishes a rule: a stated percentage of net advertising revenue for the period goes into a Dividend Pool, split by a printed formula — say half to creators pro-rata by the watch or read time they drew, half to audience members pro-rata by verified participation. Balances, revenue, and the split are posted each period.

The arithmetic. A small channel nets $2,000 a month in advertising revenue. At a 30 percent dividend, the pool is $600: $300 across creators, $300 across, say, 3,000 active audience members — ten cents each. Honest, and small. At $200,000 a month net, the same rule pays $30,000 to each side. The model's candor requirement is the point: publish the per-person number even when it is a dime. A dime stated is worth more than a fortune implied.

The contribution, named. Attention — watching, reading, returning. This model satisfies the definition only if attention is named as the contribution, in those words, wherever the Dividend is described.

The precedent. Partial analogue for the mechanic — and full precedent for the printed share. Microsoft Rewards runs attention-for-points at platform scale: points for searches, quizzes, and browser use, funded from a budget that includes search advertising revenue. Microsoft does not publish what fraction of search revenue returns as points (third-party estimates put the effective rate at roughly a tenth of a cent per point). The printed share itself is already running elsewhere: Brave's advertising program has published a 70 percent share of ad revenue to users since 2017. Attention-for-rewards is found; a published percentage of revenue is found; the model is the

combination, stated as a standard any ad-supported property can adopt.

The verdict. Workable now, thin at small scale. Any ad-supported property can adopt the rule tomorrow; the constraint is arithmetic — at typical online rates, a participant-hour of attention is worth cents, and the Dividend cannot promise more than a share of those cents. Fraud — automated accounts farming participation — is the main operating cost. Pay in cash, credit, or prizes; paying audience shares in a token with expected appreciation attaches securities flags, and the amount must be printed as varying with revenue, because it does.

Model 6 — The Data Dividend [FOUND mechanic / ORIGINAL fixed-share commitment]

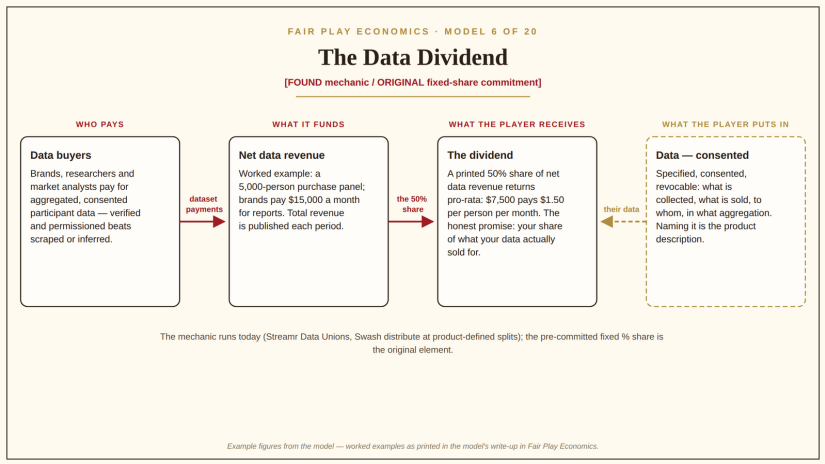


Diagram — The Data Dividend: who pays, what it funds, what the player receives, what the player puts in.

The payer and their reason. Data buyers — brands, researchers, market analysts — pay for aggregated, consented

participant data: purchase histories, panels, telemetry. Verified, permissioned data is more defensible, and often richer, than scraped or inferred data. The Dividend is the participants' printed share of that payment, paid pro-rata.

The flow. Participants opt in to specified data sharing — the specification *is* the consent: what is collected, what is sold, to whom, in what aggregation. Buyers pay the operator for datasets and reports. A printed share of net data revenue distributes pro-rata to contributing participants each period, with the total revenue and the per-person dividend both published.

The arithmetic. A 5,000-person purchase panel: brands pay $15,000 a month for reports. At a 50 percent printed share, $7,500 returns — $1.50 per person per month. At 500,000 participants and $1.5 million a month, it is the same $1.50: data revenue per head is roughly linear, so the honest promise is "your share of what your data actually sold for," which at real prices is small, and must be said. The model's claim is fairness of split, not size of payment.

The contribution, named. Data — specified, consented, revocable. Naming it is the product description.

The precedent. Partial analogue, in two halves that exist separately. Fetch sells brands aggregated, anonymized purchase-data insights, and pays participants in brand-funded points — but publishes no revenue share or dividend rate. Whether that constitutes "selling your data" is genuinely disputed: one consumer-finance analysis answers "the short answer is yes" about shopping behavior sold to companies, while Fetch's own framing is aggregated, anonymized insights with no personal identifiers or card data. Both characterizations

are reported here because the disagreement is about the definition of "your info," and a reader deserves the distinction rather than a side picked for them. Google Opinion Rewards shows the direct-payment half: researchers pay Google; respondents receive credit or cash per survey, roughly $0.10 to $2 — and the rewards expire after one year, a breakage feature the Dividend should not copy. Live data-union products distribute pooled data revenue to contributors pro-rata, at splits the product defines. The element no operator was found offering is a *pre-committed fixed* percentage of data revenue, printed in advance and held to — a rule-level commitment, recorded here at its true, modest size.

The verdict. Workable at scale. Below a few thousand participants the data is worth little and dividends are symbolic; above that, the model is a normal panel business with a fairer split. Panel data pricing runs to low dollars per participant per month — any design promising more is inventing revenue. The legal flags are the heaviest of this chapter's family: privacy-law consent and "sale/share" definitions (California's CCPA/CPRA; the GDPR for any European participants), revocation mechanics, and the rule that consent given for one dataset cannot silently stretch to another.

Model 17 — The Research Bounty Board [FOUND]

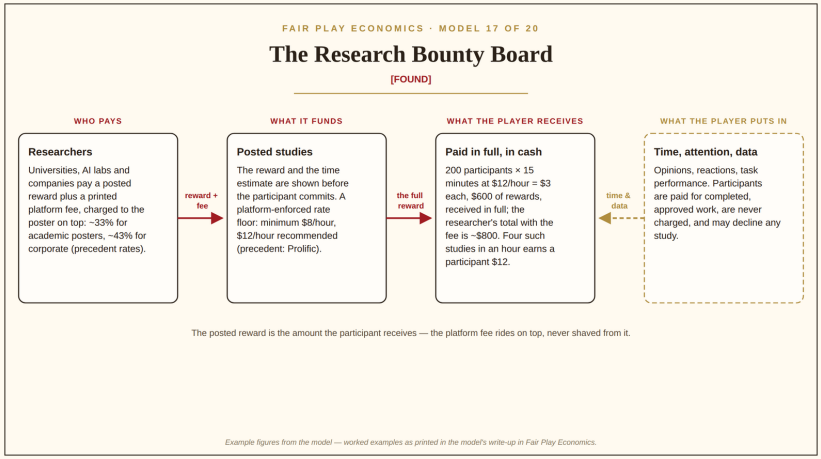


Diagram — The Research Bounty Board: who pays, what it funds, what the player receives, what the player puts in.

The payer and their reason. Researchers and institutions — universities, laboratories, companies — pay for human participation in studies and tasks, because data from verified humans is the input their work runs on. They pay a posted rate plus a printed platform fee; participants are paid for completed, approved work and are never charged.

The flow. Posters list studies with the reward and the time estimate shown *before* the participant commits. A platform-enforced minimum hourly rate floors the market. The platform fee is charged to the poster, on top of the reward — the posted reward is the amount the participant receives, not a gross figure quietly shaved. Screening is pre-matched from profiles, so unpaid screen-outs are structurally minimized, and participants may decline any study.

The arithmetic. A researcher needs 200 participants for 15 minutes at a recommended $12 an hour: $3 each, $600 of rewards. At a 33.3 percent academic platform fee, the

researcher's total is about $800, and participants collectively receive the full $600. A participant completing four such studies in an hour earns $12 — real money, paid in cash in the precedent, with no points conversion.

The contribution, named. Time, attention, and study data — opinions, reactions, task performance. The pre-commitment display of pay and time is the model's honesty mechanism and should be treated as non-optional in any adaptation.

The precedent. Full precedent. Prolific, founded at Oxford in 2014, runs the model with more than 300,000 vetted participants: an enforced minimum reward of £6/$8 per hour with £9/$12 recommended; platform fees of 33.3 percent for academic and nonprofit posters to 42.8 percent for corporate posters, charged on top of rewards; customers including Google, the National Institutes of Health, and Carnegie Mellon; payment in cash via PayPal. Prolific also demonstrates the quality side of the design: identity verification and pre-screening are what make researchers willing to pay real rates — the fee funds the trust that funds the rates.

The verdict. Workable now; cold start is the whole game. Both sides must arrive together — participants without studies leave, researchers without participants cannot recruit. The model clears only where research demand is steady; in the precedent, AI evaluation tasks are the current demand engine. Participant classification, minimum-wage analysis if an adaptation drifts toward directed work, and research-ethics norms for academic use are the flags.

Model 19 — The Idle Engine [FOUND mechanic / ORIGINAL pool-distribution form]

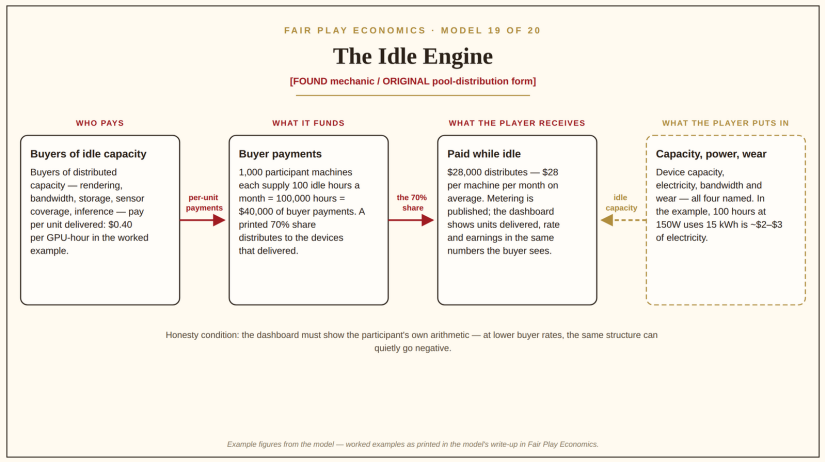


Diagram — The Idle Engine: who pays, what it funds, what the player receives, what the player puts in.

The payer and their reason. Buyers of distributed capacity — rendering, bandwidth, storage, sensor coverage, inference — pay for capacity that already exists in participants' devices and premises, because it is cheaper than building data centers for spiky demand. Participants are paid a printed share of buyer payments for capacity their devices supply while idle.

The flow. Participants install a client that supplies specified capacity only under specified conditions — idle, charging, on unmetered connections; the conditions are part of the consent. Buyers pay the operator per unit delivered. A printed share of buyer payments — say 70 percent — distributes pro-rata to the devices that delivered; metering is published, and a participant's dashboard shows units delivered, rate, and earnings in the same numbers the buyer sees.

The arithmetic. A rendering buyer pays $0.40 per GPU-hour wholesale. One thousand participant machines each supply 100 idle hours a month: 100,000 hours, $40,000 of buyer payments;

at a 70 percent participant share, $28,000 distributes — $28 per machine per month on average. Against that: a machine drawing 150 watts for 100 hours uses 15 kilowatt-hours, about $2 to $3 of electricity. The margin is real at these rates — which is exactly why the printed conditions matter. At lower buyer rates or higher power draw, the same structure can quietly go negative for the participant, and the dashboard must show the participant's own arithmetic, not just the gross.

The contribution, named. Device capacity, electricity, bandwidth, and wear — the most physically real contribution in this book. All four components must be named, and honest metering of the electricity side is a design requirement, not a nicety.

The precedent. Partial analogue. Volunteer computing — Folding@home and BOINC — proves millions will contribute idle capacity unpaid, for a mission. Bandwidth-sharing applications (Honeygain is the known example) prove the paid version exists, and also its reputational hazards: opaque buyer use of residential connections. Distributed compute networks (Render, Golem) prove buyer demand at market rates. What was not found is the pooled, game-form version: contributors' earnings pooled and redistributed by a printed percentage as a shared prize pool, run as a rewards program with published metering rather than as a crypto network or a bandwidth reseller. That distribution rule, over a well-established core, is the model's stated element.

The verdict. Fragile. Unit economics swing with buyer rates and energy prices; abuse — emulated devices, farmed capacity — attacks the same metering fairness depends on; and the

category carries a trust problem the precedent makes vivid: if participants cannot see who used their capacity and for what, the model reads as renting out the participant's connection. The design is viable only with radical metering transparency — which is why it is in this book, and why it closes the chapter.

Chapter 13 — Funded by People Who Want You Skilled

Three models in which the payer is buying attainment — someone else's. The participant's contribution in all three is study and demonstrated learning — the Bond adds a refundable lockup, named in its entry; the difference is who pays for it and why.

Model 4 — The Patron's Ladder [ORIGINAL DESIGN]

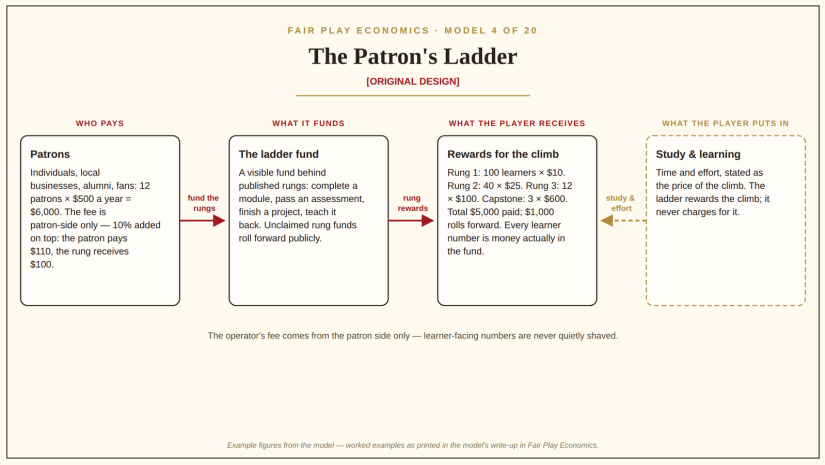


Diagram — The Patron's Ladder: who pays, what it funds, what the player receives, what the player puts in.

The payer and their reason. Patrons — individuals, local businesses, alumni, fans — fund ascending reward tiers for learners. The patron pays because the ladder is visible, named, and narratable: *"the third rung is funded by ___"* is a form of patronage as old as prizes themselves. The operator's fee comes from the patron side only — a printed percentage of patron funds, or a flat listing fee — never from learners.

The flow. A learning program publishes rungs: complete a module, pass an assessment, finish a project, teach it back. Each rung carries a reward funded by patrons into a visible ladder fund. Learners climb free. Patrons choose rungs to fund; unclaimed rung funds roll forward publicly. The fee is collected on top of patron contributions — the patron pays $110, the rung receives $100 — so the learner-facing numbers are never quietly shaved.

The arithmetic. Twelve patrons × $500 a year = $6,000, with the 10 percent fee collected as $600 on top. Rungs: 100 learners × $10, 40 × $25, 12 × $100, 3 × $600 for the capstone — $1,000 + $1,000 + $1,200 + $1,800 = $5,000, leaving $1,000 to roll forward. Every number a learner sees is money actually in the fund.

The contribution, named. Study and demonstrated learning — time and effort, stated as the price of the climb. The ladder rewards the climb; it never charges for it.

The precedent. Partial analogue. Employer-funded learning at scale (Model 18) proves institutions will prepay education; prize philanthropy — named prizes, scholarship funds — proves patrons fund visible attainment. No standing platform was found that sells *rungs of a public ladder* to patrons with a patron-side-only fee. Scholarship funds are the closest relative, but they fund *access*, not attainment rewards, and their fee structures are not printed on the ladder. The ladder as a fundable, visible structure appears original.

The verdict. Workable now; patron retention is the business. The mechanics are simple — a fund, a ladder, assessments. Completion funnels in open online learning are steep, often

under 10 percent reaching late rungs, which is favorable to fund solvency; publish expected completion rates per rung so patrons see the same math the operator sees. The real cost centers are assessment integrity and patron renewal. Rewards for attainment are income to the learner for reporting purposes; and if any rung were awarded by chance rather than attainment, sweepstakes analysis would replace skill-contest analysis — keep the rungs earned.

Working exhibit: Chess Tactics Ladder (freedough.pages.dev/chess), aboard FreeDough — twenty-six forced-mate puzzles climb in three ladders, and a miss teaches the idea.

Model 18 — The Training Purse [FOUND]

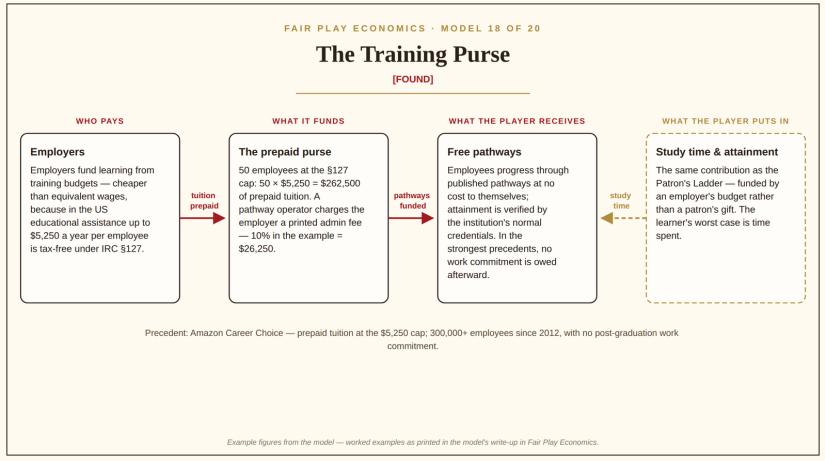


Diagram — The Training Purse: who pays, what it funds, what the player receives, what the player puts in.

The payer and their reason. Employers fund employee learning from training budgets, because skilled, retained employees are worth more than the tuition — and because, in

the United States, educational assistance up to $5,250 per employee per year is tax-free under IRC §127, which makes the purse cheaper than equivalent wages. Unused training budgets are use-it-or-lose-it; the Purse gives them a structured place to go.

The flow. The employer prepays a partner institution or a program operator. Employees progress through published pathways at no cost to themselves; attainment is verified by the institution's normal credentials. An operator running the pathways charges the employer a printed admin fee. The learner never pays — and, in the strongest precedents, owes no work commitment afterward.

The arithmetic. An employer funds 50 employees at the §127 cap: 50 × $5,250 = $262,500 of prepaid tuition. A pathway operator charging a 10 percent admin fee earns $26,250 for running pathways the employees experience as entirely free. The learner's worst case is time spent; the employer's return is measured in retention and internal mobility, which the precedents track publicly.

The contribution, named. Study time and attainment — the same contribution as the Patron's Ladder, funded by an employer's budget rather than a patron's gift. The model is kept separate from the Ladder because the payer archetype, the tax mechanics, and the precedents differ.

The precedent. Full precedent, at very large scale. Amazon Career Choice, running since 2012, prepays tuition for hourly employees at more than 400 partner schools, capped at $5,250 a year for full-time hourly associates — the §127 tax-free maximum. More than 300,000 employees have used it to pursue

new careers, with no post-graduation work commitment; Amazon has committed $2.5 billion to its broader Future Ready 2030 skills initiative. Walmart's Live Better U — no dollar cap, day-one eligibility, a tax gross-up above $5,250 — is the competing precedent and shows the design space's upper end.

The verdict. Workable now, as a business-to-business product. The precedents prove employer demand; the work is sales cycles and institutional partnerships, not mechanics. Treat the §127 cap as the natural unit size of the US purse rather than a limitation to engineer around. The flags are the written-plan and nondiscrimination requirements of §127, and the line between education benefits and taxable wages above the cap.

Model 10 — The Mastery Bond [ORIGINAL DESIGN — combination only]

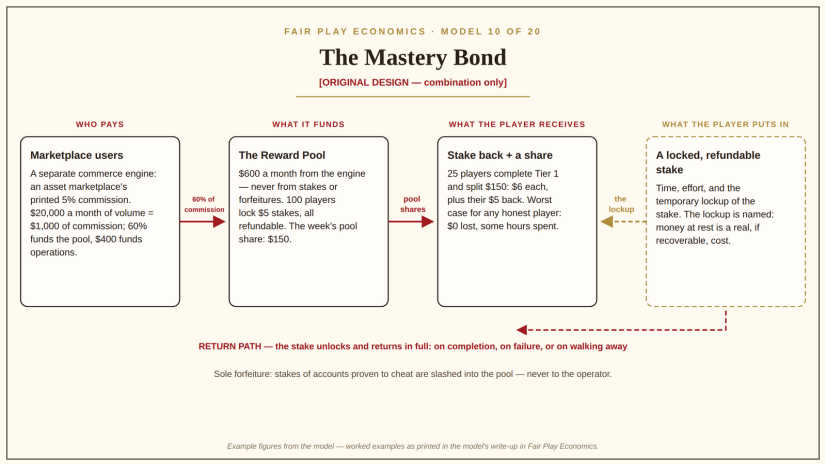


Diagram — The Mastery Bond: who pays, what it funds, what the player receives, what the player puts in.

The claim for this model is stated narrowly, because the record requires it. Refundable stakes and separately funded prize pools

each exist in prior designs — including a patent family claiming refundable entry, up to fully refundable profiles, with prize pools funded by float and third-party contributions in chance gaming. The broad claim that this model *invents* the refundable gate does not survive that record, and this book does not make it. What is claimed is a specific **combination**: a refundable stake gating a *skill*-mastery tier, returned in full on failure or withdrawal as a rule; a reward pool funded by a *separate commerce engine*, never by stakes or forfeitures; and the *sole* forfeiture being the stakes of accounts proven to have cheated, slashed into the pool and never to the operator.

The payer and their reason. The separate engine. In the worked form, an asset marketplace charges a printed 5 percent commission, of which a printed share — 60 percent — funds the Reward Pool. Marketplace users pay commissions for liquidity and discovery; skilled players are paid from the pool because their activity is what gives the marketplace's tiers and assets meaning.

The flow. To attempt a skill tier, a player locks a refundable stake. Complete the published objectives, and the stake unlocks *and* the player shares in the period's Reward Pool. Fail or walk away, and the stake unlocks after a cooldown. The gate allocates; it does not collect.

The arithmetic. A marketplace doing $20,000 a month at 5 percent commission produces $1,000; $600 goes to the pool and $400 to operations. One hundred players lock $5 stakes — $500 at rest, all of it refundable. Twenty-five complete Tier 1 and split the week's $150 pool: $6 each, plus their $5 back. The worst case for any honest player is $0 lost and some hours spent.

The contribution, named. Time, effort, and the temporary lockup of the stake. The lockup must be named, because money at rest is a real, if recoverable, cost.

The qualification, with its edges in the open. On money, the model qualifies: the stake is a lock, not a payment, and every honest participant's principal returns in full. Two edges are recorded, not hidden. First, the slashing rule means a participant *adjudged* to have cheated can lose the stake — so the design needs a published evidence standard and an appeal, or the slash becomes the operator's discretionary weapon: the prop industry's termination problem in miniature, the exact failure Chapter 4 named. Second, a stake held in a volatile asset could lose value while locked; the stake should be denominated in the same unit in which it is returned.

The precedent. Partial analogue — and a strong one. Germany's Pfand deposit system is the refundable-stake mechanic running at national scale for two decades: €0.25 on single-use containers, refunded in full on return, with return rates of about 97 to 98 percent — the highest recorded — because the stake is real enough to act on and certain to come back. Pfand proves the behavioral core: a refundable stake changes conduct at population scale without consuming anyone's money. Even the best refundable system has a residual — the unredeemed 2 percent or so is that system's breakage, and Chapter 3's question applies to it.

The verdict. Workable now, conditional on the engine and on identity. Without a second revenue engine there is no pool, and the design collapses back into a fee gate. The pool also moves with its engine: a marketplace downturn shrinks rewards in

exactly the season when participant morale matters most, and an honest operator prints that dependence rather than discovering it aloud later. Sybil farming — automated accounts completing tiers for pool shares — is the main attack; the refundable lock is itself part of the defense, because farming requires capital at rest per identity. Holding user stakes, even refundable ones, is custody: segregation and return mechanics need counsel, and slashing needs a defensible process.

Working exhibit: IQ Vault (freedough.pages.dev/iq), aboard FreeDough — fifty-two questions in four vaults; right answers earn Dough, and a miss teaches the reasoning.

Chapter 14 — Funded by Each Other, Honestly

Three models in which the community itself is the funder — of time, of surplus, and of allocation. Each is old. Each works. Each fails in a characteristic way that its chapter names.

Model 5 — The Skill Barter Bank [ORIGINAL DESIGN — the institutional bounty layer only; mechanic and platform form FOUND and old]

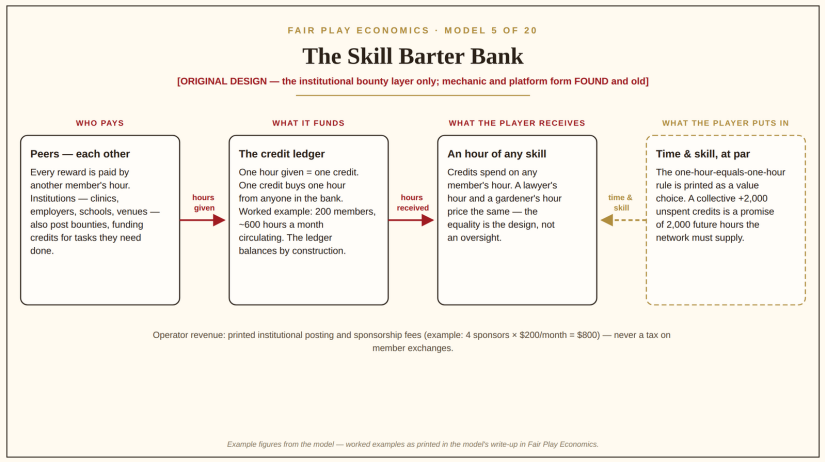


Diagram — The Skill Barter Bank: who pays, what it funds, what the player receives, what the player puts in.

The payer and their reason. Peers pay each other in kind. One hour given earns one credit; one credit buys one hour from anyone in the bank. The payer of every reward is another member's hour. The operator's revenue comes from the business layer around the bank — employers, clinics, schools, and venues pay printed fees to post bounties, tasks the institution funds in cash-to-credits or in sponsorship of the bank — never from

taxing member exchanges. That institutional bounty layer, third parties paying to post bounties denominated in the bank's own service credits, is the design's original element; the bank itself is conceded, old, and proven.

The flow. Members list offers and needs. Exchanges are logged in a shared ledger at one hour equals one credit, regardless of the market wage of the skill — the equality is the design, not an oversight. Institutions post bounties: a clinic funds credits for volunteer drivers; an employer funds a teaching series. Operator revenue is the institutional posting and sponsorship fees, printed.

The arithmetic. Two hundred members averaging three exchanges a month each: 600 hours a month circulating, with no cash moving between members. The operator's cash is, say, four institutional sponsors × $200 a month = $800, against ledger, hosting, and moderation costs. The bank's real balance sheet is hours: if members collectively bank 2,000 unspent credits, that is a promise of 2,000 future hours the network must be able to supply. Manage it like a reserve, and publish the aggregate.

The contribution, named. Time and skill, at par. The one-hour-equals-one-hour rule must be printed as the value choice it is, because it deliberately prices a lawyer's hour and a gardener's hour the same.

The precedent. Full precedent. Timebanking — Edgar Cahn's Time Dollars in the 1980s, now coordinated in the United States by TimeBanks.org and run on shared ledger software; Japan's Fureai Kippu elder-care credits predate and parallel it; Timebanking UK counted more than 25,000 members across more than 4,000 organizations by 2025. The movement's decades of operation prove the ledger works socially. They also

prove the chronic weakness: most banks survive on grants and volunteer coordination, not revenue. On taxes, this book prints the narrow statement and no more: the Internal Revenue Service's general published position is that barter income is taxable at fair market value in the year received. Private IRS determinations have held that *specific* volunteer time-dollar programs were not "barter exchanges" required to file information returns under IRC §6045 — an organization-level reporting question — expressly without opining on whether a member earns income by participating, and binding only on the organizations that requested them. The stronger claim sometimes repeated in movement literature, that time credits are tax-exempt for members, is not supported by the primary sources and does not appear in this book.

The verdict. Workable now as community infrastructure; fragile as a standalone business. The ledger balances by construction — credits are created by giving and destroyed by receiving — so the solvency risk is *service* solvency: banked credits with nothing desirable to buy. The business layer must cover real costs without taxing exchanges, which caps operator income. Say so, and the model is honest infrastructure; hide it, and the model is a grant application with a ledger attached.

Model 16 — The Patronage Pool [FOUND]

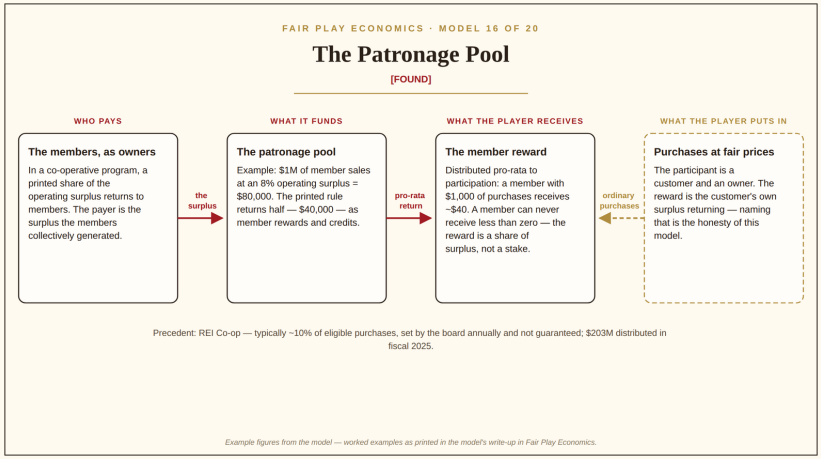


Diagram — The Patronage Pool: who pays, what it funds, what the player receives, what the player puts in.

The payer and their reason. The members themselves, in their capacity as owners. In a cooperative program or store, a printed share of the operating surplus returns to members pro-rata to their participation. The payer is the surplus the members collectively generated; the design decision is to return it by formula instead of retaining it by default.

The flow. Members transact normally, at fair prices. At period end, the surplus is computed and a printed share is distributed as member rewards or credits. The distribution rule, the surplus definition, and the limits on the board's discretion are published in advance. A member can never receive less than zero: the reward is a share of surplus, not a stake.

The arithmetic. A cooperative program with $1 million of member sales and an 8 percent operating surplus has $80,000; if the rule returns half the surplus as patronage, $40,000 flows back pro-rata — a member with $1,000 of purchases receives about $40. The numbers are modest and real, and the precedent

shows they compound into serious loyalty over decades.

The contribution, named. Ordinary purchases at fair prices. The participant is a customer and an owner, and the design must not pretend the reward is free money: it is the customer's own surplus returning. Naming that *is* the honesty of this model.

The precedent. Full precedent, at national scale, for more than 85 years. REI Co-op, founded in 1938, pays an annual Co-op Member Reward of typically about 10 percent of eligible full-price purchases, issued each March; REI distributed $203 million in member rewards in fiscal 2025 across more than 25 million members. The qualifiers are part of the precedent and are printed here because they must survive any retelling: the reward is *typically* about 10 percent and is **not guaranteed** — the board decides annually, the reward may be structured as a patronage dividend, a promotional reward, or a mix, and it is issued as merchandise credit that itself expires. Patronage dividends of $10 or more are reported on Form 1099-PATR.

The verdict. Workable now. The cooperative form is mature law and mature practice. The pool can only distribute surplus that exists — years without surplus pay nothing — and the printed rule must say so in advance; REI's "typically, not guaranteed" language is the honest template. The flags are cooperative tax treatment, which rewards proper structure, and a bright line: membership must never be marketed as an investment.

Model 20 — The Quadratic Commons [FOUND mechanic and prize-allocation practice; this book's contribution is the season framing for general programs]

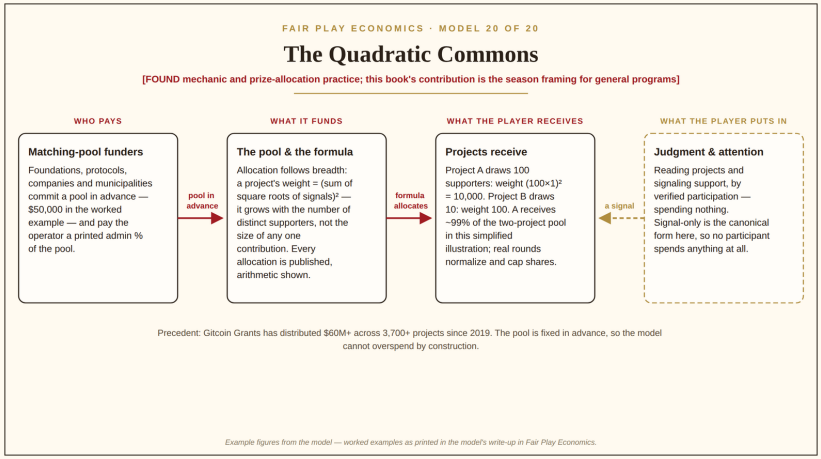


Diagram — The Quadratic Commons: who pays, what it funds, what the player receives, what the player puts in.

The payer and their reason. Matching-pool funders — foundations, protocols, companies, municipalities — commit a pool in advance. Community members then direct the pool by breadth-weighted signals: allocation follows the *number* of distinct supporters, not the size of any one contribution, under the quadratic rule — a project's weight grows with the square of the sum of the square roots of its signals. Funders pay because the mechanism discovers what a community actually values, with receipts.

The flow. A season opens with a printed pool and a printed formula. Participants signal support for projects — in the canonical form used here, by verified participation, an account in good standing casting weighted votes, spending nothing. At season close, the formula allocates the pool, and every allocation is published with its arithmetic shown. The operator is paid a printed admin percentage of the pool, by the funders.

The arithmetic. A $50,000 pool across five projects. Project A draws 100 supporters; Project B draws 10. Under the signal-only quadratic weight, A's weight is $(100 \times 1)^2 = 10{,}000$ against B's $(10 \times 1)^2 = 100$ — breadth dominates by design; real rounds normalize across all projects and cap per-project shares, but the direction is the design's whole point. The precedent's canonical illustration makes the same point in donation form: 100 donors giving $10 each attract roughly ten times the matching of 10 donors giving $100 each, on the same $1,000 of direct contributions.

The contribution, named. Judgment and attention — reading projects and signaling. In donation-based rounds, participants may also *give* money; a gift is a voluntary transfer, not a loss mechanism, and the distinction is recorded here rather than blurred — but program designs in this book default to signal-only, so that no participant spends anything at all and the qualification needs no footnote.

The precedent. Full precedent for the mechanic. Gitcoin Grants, running since 2019, has distributed more than $60 million across more than 3,700 projects through quadratic funding rounds, with matching pools from foundations and protocol treasuries — and with industrial-strength identity verification, because breadth-based allocation invites fake crowds. Municipal pilots — Split's triple-match citizen funding among them — show the form leaving its original habitat. The mechanic and its prize-allocation practice are found; this book's contribution is the *season framing*: a recurring, published, prize-like allocation season run for a program's own community projects.

The verdict. Workable now. The formula and the anti-sybil tooling both exist and are battle-tested. The pool is fixed in advance, so the model cannot overspend by construction; its failure mode is not insolvency but **capture** — coordinated blocs steering allocations, and manufactured identities inflating breadth. Identity cost is therefore a load-bearing part of the design, not an accessory, and a Commons that skimps on it will be governed by its most organized minority. Charitable-solicitation rules apply if donations are accepted, and allocating other people's money carries duties of care even outside formal charity law.

Chapter 15 — Funded by Savings and Surplus

Five models in which the payer pays out of money that exists only because the system worked: a sister business's sales, an avoided cost, an insurer's prevented claims, a funder's purchased outcome — and, in the closing model, breakage itself, rerouted.

Model 7 — The Cross-Subsidy Atelier [FOUND mechanic / ORIGINAL destination]

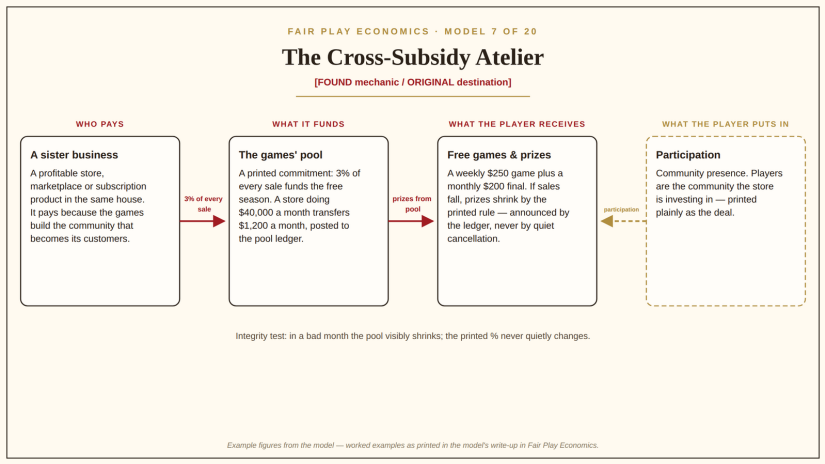


Diagram — The Cross-Subsidy Atelier: who pays, what it funds, what the player receives, what the player puts in.

The payer and their reason. A profitable sister business — a store, a marketplace, a subscription product — prints a commitment: a stated percentage of every sale, or of gross margin, funds free games and rewards run by the same house. The business pays because the games build the community that becomes its customers, and because a printed percentage converts marketing spend into a public, checkable promise

instead of advertising fog. Printed-percentage funding commitments are old and famous — Patagonia's 1% for the Planet has run since 1985. The destination is the model's element: the percentage funding free games and rewards for the seller's *own community*.

The flow. The rule is stated where customers can see it — "3 percent of every order funds the free season." Each period, the transfer is computed from published sales figures and posted to the games' pool ledger. Games are free to enter; prizes come only from the transferred pool. If sales fall, prizes shrink — by the printed rule, announced by the ledger, never by quiet cancellation.

The arithmetic. A store doing $40,000 a month at a 3 percent-of-sales commitment transfers $1,200 a month: a weekly $250 game plus a monthly $200 final. A store doing $4 million a month transfers $120,000. The model's integrity test is the bad month: the pool must visibly shrink rather than the percentage quietly changing.

The contribution, named. Participation and community presence. Players should be told plainly that they are the community the store is investing in. That is the deal, and it is a fair one when printed.

The precedent. Full precedent for the mechanic; partial for the printed-percentage form. McDonald's Monopoly, running since 1987, is the canonical brand-funded prize game — prizes funded from the marketing budgets of the restaurant business and its supplier partners, entry attached to purchases with a free mail-in entry route. Its history also supplies Chapter 6's cautionary case: the funding was honest and the pipeline was not, for over a

decade. Many modern analogues exist — grocery and fuel rewards funded by manufacturer trade spend — and few print the subsidy percentage.

The verdict. Workable now, conditional on the engine — the model is exactly as strong as the sister business. Size the commitment against *margin*, not hope: 3 percent of sales at a 10-percent-margin business is 30 percent of profit, which will not survive the first bad quarter. Print the basis — sales or margin — with the percentage. If game entry is tied to purchases, a free entry route is required.

Model 12 — The Shared-Savings Split [FOUND mechanic, single-domain / ORIGINAL cross-domain synthesis (framing)]

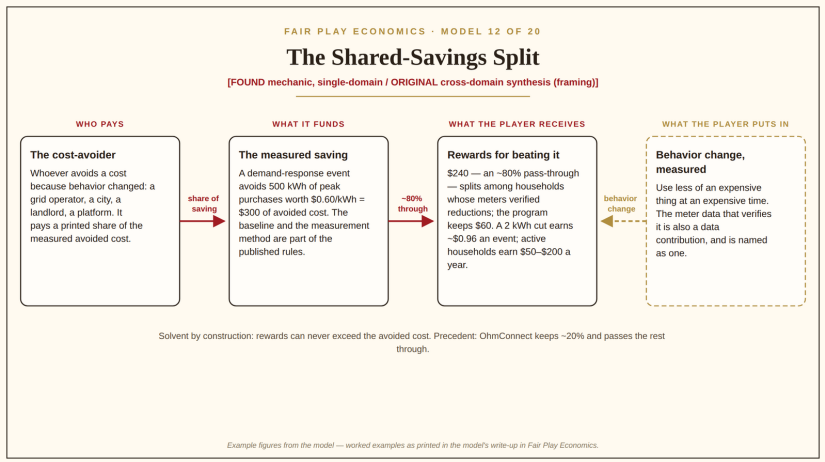


Diagram — The Shared-Savings Split: who pays, what it funds, what the player receives, what the player puts in.

The payer and their reason. Whoever avoids a cost because participants changed behavior: a grid operator avoiding peaker-plant runs, a city avoiding landfill tonnage, a landlord

avoiding vacancy, a platform avoiding support burden. The payer shares the *measured avoided cost* — money that exists only because the saving happened. Gainsharing within a single domain is established, patented practice; the model's contribution is naming the pattern across domains, a synthesis offered as a framing, not an invention claim.

The flow. A baseline is published — say, a household's forecast peak-hour energy use. Participants who verifiably beat the baseline during called events earn rewards from a pool funded by the payer at a printed share of the measured saving; the working example passes about 80 percent through and keeps about 20 percent as the program margin. The measurement method and the baseline formula are part of the published rules — the baseline *is* the game's integrity surface.

The arithmetic. A demand-response event avoids 500 kilowatt-hours of peak purchases worth $0.60 per kilowatt-hour to the grid operator: $300 of avoided cost. At an 80 percent pass-through, $240 splits among the households whose meters verified reductions, and the program keeps $60. A household that cut 2 kilowatt-hours earns about $0.96 for the event; active households in the working example earn $50 to $200 a year. Nobody paid anything to be in the pool.

The contribution, named. Behavior change, measured: use less of an expensive thing at an expensive time. The measurement — smart-meter or equivalent data — is also a data contribution, and is named as one.

The precedent. Full precedent in one domain. OhmConnect (now Ohm Energy) pays households in California, New York, and Texas to reduce electricity use during grid-stress events:

energy authorities and utilities pay for aggregated reductions because avoided peaker generation is cheaper than generation. OhmConnect has historically kept about 20 percent of the payment and passed the rest to users as cash or charity donations; typical user earnings run $50 to $150 a year, by its co-founder's account. The generalization — the same split applied to any measurable avoided cost, with split and baseline printed — is the step beyond the precedent.

The verdict. Workable now, wherever measurement exists. The model lives or dies on baseline integrity: if participants can inflate the baseline, they manufacture "savings." Rewards can never exceed the avoided cost, so the design is self-limiting and solvent by construction; its ceiling is the size of the inefficiency it attacks. Utility-market participation is regulated — the payer relationship, not the player side, carries that weight.

Model 13 — The Non-Event Dividend [FOUND]

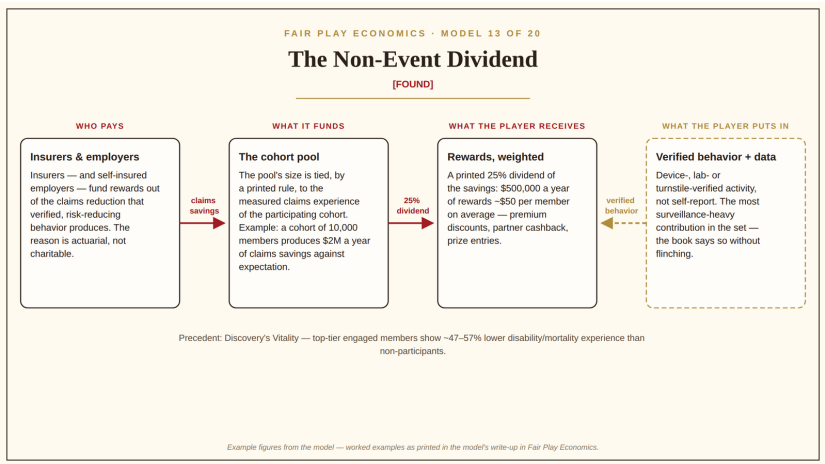


Diagram — The Non-Event Dividend: who pays, what it funds, what the player receives, what the player puts in.

The payer and their reason. Insurers — and self-insured employers — fund rewards for verified risk-reducing behavior: activity, health checks, safe driving. The payer's reason is actuarial, not charitable: the reward costs less than the claims the behavior prevents. The dividend is paid, in effect, for the event that did not happen.

The flow. Participants opt into verified behavior tracking — the named data contribution is device-, turnstile-, or lab-verified activity, not self-report. Behavior earns points and status; status converts to rewards — premium discounts, partner cashback, prize entries — funded from a pool whose size is tied, by a printed rule, to the measured claims experience of the participating cohort. The game form this book proposes is the cohort challenge: the whole cohort's verified activity this quarter funds next quarter's reward table, so the funding loop is visible to the people generating it.

The arithmetic. An insurer's engaged cohort shows materially lower claims — in the precedent, top-tier members show roughly 47 to 57 percent lower disability and mortality experience than non-participants. If a cohort of 10,000 members produces $2 million a year of claims savings against expectation, a printed 25 percent dividend is $500,000 a year of rewards — $50 per member on average, weighted by verified behavior. The program operator is paid by the insurer from the same savings, at a printed rate.

The contribution, named. Verified behavior, and the data that verifies it. This is the most surveillance-heavy contribution in the set, and this book says so without flinching, here as in Chapter 8: the participant trades continuous proof of behavior

for rewards funded by their own risk reduction.

The precedent. Full precedent for the funding loop. Discovery's Vitality, running in South Africa since 1997 and now in more than 40 markets covering about 49 million lives, pays members points for verified healthy behavior, with rewards including up to 75 percent off gym memberships and 25 percent cashback on healthy food. Discovery Life paid R2.4 billion in shared-value rewards in 2024 — R1.4 billion in PayBacks and R1.0 billion in cash conversions — funded by the actuarial loop: engaged members' lower claims. The precedent is reported with its critique, because both belong in the record: academic critics argue the model also commodifies health data and shifts risk responsibility onto individuals. The funding loop is proven; its social terms are contested.

The verdict. Workable at scale, or with a partner. A standalone publisher cannot be its own insurer; the model requires an insurer or employer partner, or a narrow self-funded version — a company funding its own workforce's dividend. Rewards must trail *measured* savings by at least a year of claims data; paying ahead of the actuarial evidence is how such programs quietly die. Insurance rebating and inducement rules vary by state, and behavior-based pricing deserves anti-discrimination review alongside the privacy analysis.

Working exhibit: The Calm Watch (freedough.pages.dev/calm), aboard FreeDough — the boat's mutual collects the Café's daily premium, pays a dividend to everyone who stood the watch on the days the bad day does not happen, and pays repairs out of the same pool when it does.

Model 14 — The Outcome Buyer's Table [ORIGINAL framing — printed per-outcome price table paying producers directly; funding family FOUND]

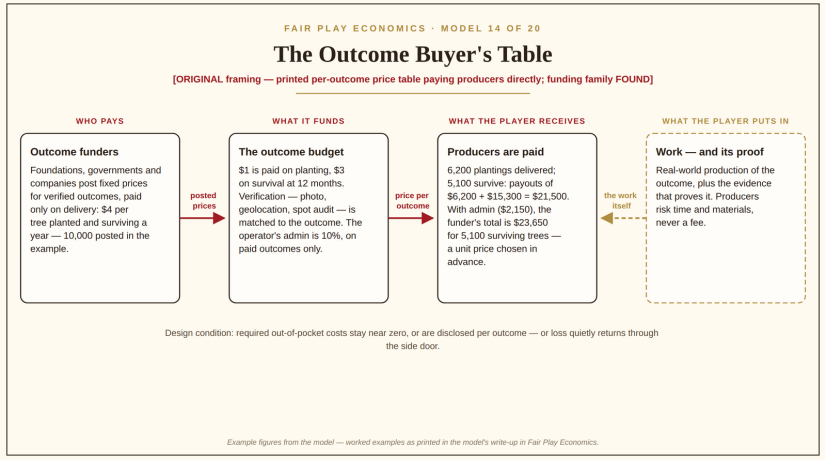


Diagram — The Outcome Buyer's Table: who pays, what it funds, what the player receives, what the player puts in.

The payer and their reason. Outcome funders — foundations, governments, companies with commitments — post fixed prices for verified outcomes: a stated amount per tree planted and surviving a year, per learner certified, per kilogram of electronic waste collected, per home weatherized. The funder pays because it wants the outcome itself, priced in advance, paid only on delivery. Results-based financing is the found funding family; the model's framing is the element claimed here: funders posting *printed per-outcome prices* whose purchases pay rewards *directly to the participants who produce the verified outcomes*.

The flow. The Table publishes outcome prices and verification standards. Anyone may produce an outcome and claim its price; verification matched to the outcome's cost of proof — photograph plus geolocation plus spot audit, certification

records, weighbridge receipts — releases payment. The operator charges the funder a printed admin percentage on paid outcomes only. Unsold outcome budget returns to the funder or rolls to the next season, printed at posting.

The arithmetic. A foundation posts 10,000 verified tree-plantings at $4, with survival checked at twelve months: $1 on planting, $3 on survival. Producers deliver 6,200 plantings; 5,100 survive. Payouts are $6,200 + $15,300 = $21,500; admin at 10 percent is $2,150. The funder's total cost is $23,650 for 5,100 surviving trees — a unit price it chose before spending anything. Producers risked time and materials, never a fee.

The contribution, named. Work — real-world production of the outcome, plus the evidence that proves it. One honesty condition is part of the design: producers may spend their own money on materials pursuing an outcome, so required out-of-pocket costs must be kept near zero, or disclosed per outcome — otherwise the model quietly reintroduces loss through the side door.

The precedent. Partial analogue. Results-based financing and development impact bonds — the Educate Girls development impact bond of 2015 to 2018 is the standard example — pay for verified outcomes, but in an impact bond the *investors* front the working capital and can lose it. The Table removes the investor layer and pays producers directly; that structural difference is why producers here cannot lose money. Pay-for-success contracts in US social policy are cousins. No standing open Table — any producer, posted prices, direct claim — was found.

The verdict. Workable now, for cheaply verifiable outcomes. The binding constraint is verification cost per outcome: if proof

costs more than about 15 percent of the outcome price, the Table eats itself. Digital-native outcomes — certifications, merged code, labeled datasets — verify cheapest; physical outcomes need audit sampling. Fraud pressure scales with price, and Chapter 6's lesson applies to verifiers as much as to prizes. The flags are contractor and employment classification for producers working at volume, and charitable-funds handling where foundation money is the pool.

Working exhibit: Predictions Desk (freedough.pages.dev/predictions), aboard FreeDough — the Desk takes calls on sealed readings, returns every stake in full, and pays a printed prize for a right call.

Model 11 — Reverse-Breakage Commons [ORIGINAL DESIGN]

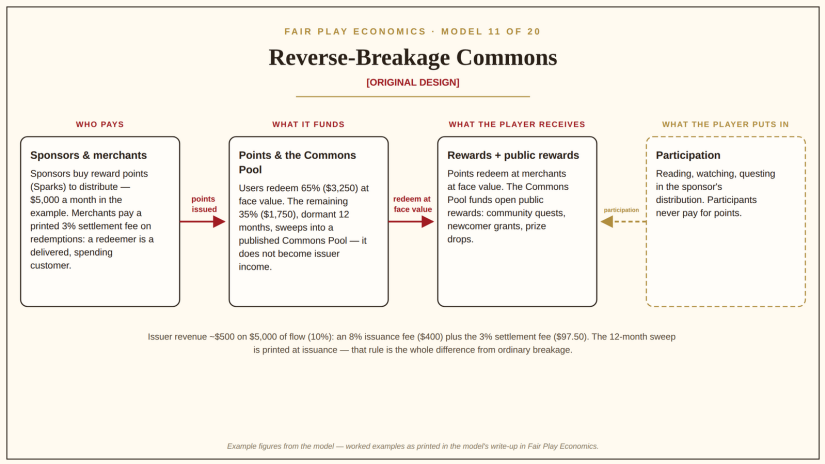


Diagram — Reverse-Breakage Commons: who pays, what it funds, what the player receives, what the player puts in.

The last model in this chapter closes the loop Part I opened. Chapter 3 described breakage: value that evaporates becomes

the issuer's income. This design keeps the evaporation — people will forget points in any system ever built — and changes who receives it.

The payer and their reason. Sponsors buy reward points to distribute as marketing; merchants pay a printed settlement fee on redemptions, because a redeeming customer is a delivered, spending customer. The issuer is paid for issuance and plumbing — never for value that quietly evaporates.

The flow. Points are issued against sponsor funding and redeemed at merchants at face value. Points dormant for twelve months are **not** cancelled into issuer income: they sweep into a published Commons Pool that funds open public rewards — community quests, newcomer grants, prize drops. Issuer revenue is a printed issuance fee plus a printed settlement fee on redeemed value.

The arithmetic. A sponsor buys $5,000 of points monthly. The issuance fee at 8 percent is $400. Users redeem 65 percent — $3,250 — producing settlement fees of $97.50 at 3 percent. The remaining 35 percent, $1,750, would in a conventional program become breakage income; here it is a liability that migrates to the Commons Pool after dormancy. Issuer revenue is about $500 on $5,000 of flow — 10 percent. That is less than the breakage model's potential capture, and the whole of it prints on one page.

The contribution, named. Participation in the sponsor's distribution — reading, watching, questing. Participants never pay for points.

The qualification, stated plainly. The model qualifies: participants lose nothing they paid in, because they paid nothing

in. But the dormant value that sweeps to the Commons *was* the participant's earned reward, and the design must say so: after twelve months dormant, unredeemed points become community property **by a rule printed at issuance**, not by a term discovered later. That printed-at-issuance rule is the entire difference between this design and ordinary breakage.

The precedent. None found. Loyalty breakage is universal — the ranges of Chapter 3: 10 to 19 percent of gift-card value, roughly 15 to 40 percent of loyalty value — and no issuer was found that routes its own breakage to a visible community pool instead of income. The inversion appears genuinely unused, plausibly because breakage *is* the margin for many issuers.

The verdict. Fragile — and the fragility has a name, stated in Chapter 7 and repeated here because it governs this model: **escheat**. Dormant balances may be unclaimed property owed to the owner or to the state, not to any pool an issuer designs. Whether earned points are "property" for unclaimed-property purposes varies by state and by how the points were acquired — purchased value is treated more strictly than promotional value, and promotional cards sit outside the federal gift-card protections, which cuts both ways: fewer consumer-protection constraints, and no federal safe harbor either. The design likely works in some jurisdictions with points structured as promotional, non-purchased value and a dormancy rule disclosed at issuance — but this is the set's clearest "lawyer before launch" item, and no honest version of the model ships without that review. Two secondary fragilities ride along: the merchant-network bootstrap, and Commons governance — who designs the quests the Pool funds is a real institution to run, with real capture risks of the kind Chapter 14 named, and it deserves

the same printed-rule treatment as the money.

Working exhibit: Bingo Commons (freedough.pages.dev/bingo), aboard FreeDough — cards are free, the prize tables are printed, and the caller seals the full draw order before the first ball.

Closing Case — The Saturday League

Part III ends not with a twenty-first model but with three of its models running together in one small, picturable program — because the lesson of a stack is one no standalone model can teach. The League is an **original composition**: an arrangement of Model 1 (the Merchant's Table, [FOUND]), Model 2 (the Endowment Season, [ORIGINAL game-season arrangement; funding structure FOUND]), and Model 5 (the Skill Barter Bank, [ORIGINAL DESIGN — the institutional bounty layer only; mechanic and platform form FOUND and old]). The components keep their own labels; the composition is the original element.

The program. A free-entry, twelve-week skill season — a local league, agnostic about the skill: a pub-quiz ladder, a chess or racket league, a makers' challenge series — with 60 players. Nobody pays to enter, at any point, for anything. Three unrelated payers fund three different layers, each buying something different.

Engine 1 — the Merchant's Table funds the weekly prizes. Ten local shops each pay $25 a week into the league pool against verified player visits: the league's weekly route brings players past and into the shops, and visits are counted by a published verification method. That is $250 a week, $3,000 across the season, split by printed rule: $200 a week in weekly prizes ($2,400 across the season) and $50 a week banked to the championship ($600). What the merchants are buying is verified foot traffic — 60 local players, routed weekly, at $25 a week flat

— priced below what the same visits cost through ordinary local advertising, and paid only against a count they can audit.

Engine 2 — the Endowment Season funds the championship. A founding patron endows $20,000, once. Invested conservatively under Model 2's spending rule, a yield of about 4 percent — roughly $800 a year — funds the championship prize in perpetuity; the principal is never spent. The season-one championship pool is the endowment's $800 distribution plus the $600 banked from Engine 1: $1,400. From season two onward, the endowment's distribution arrives whether or not any given season's merchants renew. What the patron is buying is perpetuity and a name — a championship that carries the patron's name or cause every year, indefinitely, at a one-time cost: the Nobel purchase at neighborhood scale. The design scales down without changing shape — a $5,000 endowment yields about $200 a year, a smaller championship under the same rules — and $20,000 is a worked figure, not a threshold.

Engine 3 — the Skill Barter Bank runs the coaching layer. Veteran players teach one session a week, earning barter credits at the Bank's printed par — one hour given, one credit — and newcomers spend credits on coaching. No money moves in this layer at all. The veterans receive standing, credits spendable on each other's skills, and a league whose newcomers improve: their future competition and their community. The league receives a retention and onboarding layer that costs the cash engines nothing.

The contribution, named. Presence and foot traffic (Engine 1's value), participation (Engine 2's audience), and time — playing, and for veterans, teaching (Engine 3). Money is never among the

contributions, and the program's materials say so in those words.

Failure isolation — the point of the stack. Each engine can fail without bankrupting the others, and no failure creates a debt to participants, because no participant ever paid in. If three of the ten merchants drop out in week six, the weekly pool falls to $175 by the printed rule — $140 in prizes, $35 banked; the championship, funded by the endowment and banked funds already collected, is untouched, and no player is owed a refund of an entry fee that never existed. If the endowment has a bad year, the spending rule shrinks the distribution: the championship gets smaller, nothing is cancelled retroactively, and no promise breaks, because the printed rule promised a *share of the distribution*, never a fixed amount the endowment might not produce. If the barter bank thins out, coaching gets scarcer — a service-solvency problem, in Model 5's terms, not a cash problem — and the prize engines are entirely unaffected. If the whole league folds, the endowment, governed by its own instrument, redirects or persists for a successor season; merchant obligations end with the last verified week; barter credits lapse as the in-kind promises of a community that no longer meets — a real disappointment, and a loss of exactly $0 to every participant.

The verdict. Workable now. Every component is individually rated workable now in this set, and the stack adds no new mechanism — only coordination: one printed rule sheet covering all three engines, one verification count the merchants accept, one endowment instrument. The founding act is the honest bottleneck: the $20,000 endowment must exist before Engine 2 does. The teaching of the case is diversification. Three payers with three unrelated reasons — foot traffic, perpetuity,

community — is what Chapter 5's rule, *arithmetic that survives a bad month*, looks like when it is built rather than stated. This is where the twenty models stop being a catalog and become a kit.

Conclusion — A Five-Minute Audit for Any Program

Chapter 9 stated the standard from two chairs. This page is the tool form of that standard — six questions, applicable to any game, reward program, or token system, before entering it or building it. Five minutes is enough, because every question is answered by documents the program itself publishes, or by their absence.

1. **Who pays?** Name the payer — a person, a company, a fund. If the answer is a fog ("the ecosystem," "the community," "the pool"), the answer is the participants.
2. **Why is paying worth it to them?** State the payer's reason in one sentence a skeptic would accept: a delivered customer, a perpetuity, an avoided cost, a delivered outcome. A payer without a reason is a subsidy, and subsidies end.
3. **Is the rate printed?** The fee, the share, the split, the expiry rule — published before entry, where entrants actually read. A rate that exists only in a filing or a support answer is not a printed rate.
4. **Who controls the prize pipeline?** Who holds the prizes, who can reach them, and who besides the operator has checked? Fair funding with an unaudited pipeline is not fair; Chapter 6 is the record of why.
5. **What do I contribute, and is it named?** Money, attention, time, data, behavior, lockup — stated on the label, in plain words, with its weight honestly described. Zero-loss is a

claim about money, not about costlessness; a program that hides its contribution has priced it at whatever it can take.

6. **What happens to value nobody claims?** Breakage, dormancy, unspent pools, shutdown. The destination of unclaimed value is the single most revealing line in any program's design — it is where the participant-funded systems of Part I collect their quietest revenue, and where the honest designs of Part III print their answer at issuance.

A program that answers all six, in writing, before you commit, has earned the next hour of your attention. A program that cannot answer the first has answered it.

One design is owed a closing note, because its exclusion is this book's definition keeping itself. The research behind Part III produced one further design, the **Glass Pool**: an entry-fee prize pool split by a printed, unchangeable formula — 70 percent to prizes, 20 percent to a player-owned reserve, 10 percent to operations — with the reserve publicly reported, balance and every movement; topping up thin weeks to the promised prize table; capped at a printed multiple, with the overflow above the cap returned to players; untouchable by the operator for any reason; and, on shutdown, distributed to recent players by a printed formula. The design at working scale: a weekly arena takes 200 entries at $5 — $1,000 in — and pays $700 in prizes, banks $200 in the reserve, and keeps $100 for operations. With steady play and no drawdowns, the reserve reaches its $2,800 cap in week 14; from week 15 the $200 stream overflows to players, and the expected prize value per entry rises from $3.50 to $4.50 on a $5 ticket against a known $0.50 operator take. A thin week — 80 entries, $400 in — still pays the full $700 table: $280 comes from that week's split and the reserve covers the

$420 gap. The operator's annual take at this scale is about $5,200: honest, modest, and exactly as advertised, because the design deliberately caps operator upside in exchange for zero outcome risk. It is the most transparent entry-fee pool the research found or built. It is not in Part III, because its prizes are funded from entry fees, and an entrant who does not win ends with less money than they started with. That is exactly what this book's definition excludes.

The exclusion stands on the definition — and the design's originality, examined, is narrower than it first appeared, which is also worth printing. The reserve core is already running in published lottery rules: Powerball's MUSL Prize Reserve and the EuroMillions Reserve Fund both print the split, segregate the reserve, top up deficiencies, cap the fund, and overflow the excess to players. What no product was found doing is the design's transparency-and-exit layer: reporting the reserve's balance and every movement publicly, and distributing it to recent players by printed formula on shutdown. A book called *Nobody Has to Lose* keeps its definition even when it costs a good design — and says, precisely, what the good design still had that the world has not used.

The paintings at the front of this book and at the head of its parts come from FreeDough, the arcade where these designs run as games anyone can play. That is the relationship, stated once and plainly: this book sets out the designs, and FreeDough is the same designs running live — the arcade is the book's working proof, and the book is the arcade's workings, shown. A reader who wants to test any claim in Part III against a running system can walk into the arcade and read the payer's name on the door, because the arcade prints it there — which is, in the end, the

whole request this book makes of every program it describes.

Ask who pays. Keep asking until the answer is a name and a reason. Then check the arithmetic, read the contribution on the label, and look at who holds the prizes. Systems in which nobody has to lose are not rare because they are impossible. They are rare because they are built on purpose — and building on purpose begins with the question.

That is the argument. What follows is a theory — the author's, labeled as one.

Games people love, teaching things people need to learn, gather the one resource every money game runs on: attention, given willingly and returned to daily. The models in this book say that attention never has to be farmed out of the player's pocket. Sponsors, advertisers, merchants, outcome buyers, and buyers of consented, openly brokered game data will fund the prizes — because a delivered customer, a verified outcome, and an honest dataset are each worth a printed price to the party that pays it.

If the payers are third parties and the rates are printed, play can be free and the arithmetic fair, and the system can still run cash-flow positive on both sides of the counter: the player risks nothing and can come out ahead; the builder earns from the funding, not from the losses. Run that way, the ecosystem turns green — self-funding and growing — and it clears an honest path through problems the extractive designs cannot solve, because fraud and quiet executive manipulation live in hidden money flows, and this design has none: the payer is named, the pipeline is printed, and the audit takes five minutes.

The theory stands or falls on a measurable claim: that third-party funding, at printed rates, covers prizes and operations at the scale popularity provides. The arcade bound into this book is the first small test of it, running.

A request from the author: use these models. Make gaming educational, free, and fair — and let it pay, honestly, the people who play it and the people who build it.

The Working Exhibits

As of November 2026, these exhibits are open aboard FreeDough, the working arcade at freedough.pages.dev. Each entry names the model it exhibits, so a reader can test the printed design against the running version.

- **The Café (freedough.pages.dev/#arcade) — Model 1, The Merchant's Table.** The café and shops fund the games because the games bring the crowd — the explanation of who pays, printed on the door.
- **Chess Tactics Ladder (freedough.pages.dev/chess) — Model 4, The Patron's Ladder.** Twenty-six forced-mate puzzles climb in three ladders. Attempts are free, and a miss teaches the idea.
- **The Board (freedough.pages.dev/board) — Model 8, Sponsor-the-Board.** Sponsors post bounties across six lanes — United States trivia, world and history, science, math, spelling, and crypto knowledge — and players claim them by answering.
- **IQ Vault (freedough.pages.dev/iq) — Model 10, The Mastery Bond.** Fifty-two questions in four vaults — number, word and letter, pattern, and logic. Right answers earn Dough; a miss teaches the reasoning.
- **Bingo Commons (freedough.pages.dev/bingo) — Model 11, Reverse-Breakage Commons.** Free cards in three forms, with printed prize tables. The caller seals the full draw order before the first ball and reveals it when the card ends.

- **The Calm Watch (freedough.pages.dev/calm) — Model 13, The Non-Event Dividend.** The boat's mutual collects the Café's daily premium into a pool; on the days the bad day does not happen, a dividend from the pool is split among everyone who stood the watch, and on the days it does, the pool pays repairs.
- **Predictions Desk (freedough.pages.dev/predictions) — Model 14, The Outcome Buyer's Table.** Three desks take calls on sealed daily readings. Every stake returns in full at settlement, and a right call earns a printed prize.
- **The Slot Fleet (freedough.pages.dev/fleet), including The Honest One (freedough.pages.dev/slots) — composite exhibit.** Fifty-one machines print their reels, paytables, odds, and average return. The Café funds the pools and every stake returns; the flagship's own plate names the composite: the Merchant-Funded Machine, fed by the Reverse-Breakage Commons, with the stake held as a bond.

The remaining models in Part III are not yet aboard. They are the building list for later waves of the arcade; as each new door opens, this page changes in the next edition.

Support

Support this work with Bitcoin: bc1qdtqldz6enu5gw4hdq8kq0wq8thykp085ka439j

About the Art

The four paintings in this book come from FreeDough, the working arcade built on the models in Part III, where they hang as the top shelf of its prize counter — the shelf reserved for the rewards a player earns only over a long season of play. They are reproduced here as plates because the tie between the book and the arcade is the point: this book sets out the designs, and FreeDough is the same designs running live, where anyone can play them. The arcade is the book's working proof, and the book is the arcade's workings, shown.

The four plates, in order: The Riverboat Herself (frontispiece); The First Baker, opening Part I; Your Named Star, opening Part II; Every Port of Call, opening Part III. Paintings by Douglas John Brownfield.

The twenty mechanics diagrams in Part III were drawn for this book; each carries its model's own worked figures, as printed in its entry.

About the Author

Douglas J. Brownfield is an independent builder and artist based in America. He builds self-custodial software tools — programs designed so the people who use them keep hold of their own work — and he writes about what it is actually like to work with AI day after day: where it shines, where it guesses, and where it has to be watched.

This book began with a question the author kept asking of every money game he met: who pays? Following it through prop-firm challenges, loyalty points, and tournament pools produced the research in Part I. The twenty designs in Part III are the answer, built and then run: FreeDough, the arcade where the models operate as games anyone can play, is this book's working proof, and its paintings hang in this book as plates.

The author's other books stand on the same shelf rule: one subject, primary sources, no drama framing — the facts at the strength the evidence supports, and the open questions left open and labelled.

Also by Douglas J. Brownfield

Words to Prompt AI Well With — a beginner-friendly field guide to asking AI better questions and getting better answers.

Before Q-Day — how to prepare your Bitcoin and crypto for the quantum change: what's settled, what's coming, and what to do in the meantime.

What Kind of AI Is That? — the types of AI models, what each kind does, and how to tell them apart.

Even Further — Hubble and the Nancy Grace Roman Space Telescope: two telescopes built around the same 2.4-metre mirror, thirty-six years apart — what they're capturing, and what comes next.